

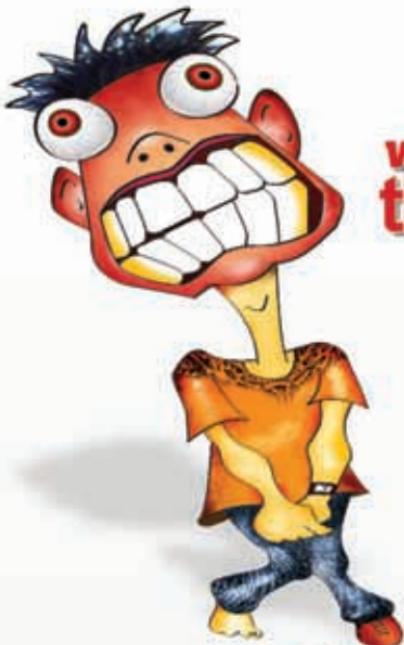
digit & FastTrack

YOUR HANDY GUIDE TO EVERYDAY TECHNOLOGY

to ARTIFICIAL INTELLIGENCE



- Ancient AI?
- Radical thinkers
- Types of AI and tests
- Intelligence - The human mind and creativity
- Fictional AI
- Intelligence today
- ChatterBots
- IBM Watson
- AI and the future
- Looking at AI under a new light



[www.
thinkdigit/forum](http://www.thinkdigit.com/forum)

Join the forum to
express your views
and resolve your
differences in a more
civilised way.

thinkdigit
FORUM

Post your queries
and get instant
answers to all
your technology
related questions



one of the most active online technology forum
not only in India but world-wide

**JOIN
NOW**



www.thinkdigit.com



ARTIFICIAL INTELLIGENCE

powered by



CHAPTERS

ARTIFICIAL INTELLIGENCE

DECEMBER 2011

06
PAGE

Ancient AI?

Myths and tales which inspired reality – right from Adam (of Adam & Eve fame) to Pandora and Frankenstein's Monster. They were just some of the earliest references to Mechanical Intelligence.

11
PAGE

Radical thinkers

Alan Turing, John McCarthy, Marvin Minsky, Arthur C Clarke and Isaac Asimov...read on to know why these names are as revered as they are.

16
PAGE

Types of AI and tests

Did you know that there are two types of Artificial Intelligence? Once you understand the difference between these, learn about the Turing test used to create intelligent agents and why it's so very important.

26
PAGE

Intelligence – the human mind and creativity

How does human logic and creativity apply to artificial intelligence? Why is creativity so important in developing AI anyways? Can an AI agent have fun at your expense? Find out in this section.

CREDITS
The people behind this book

EDITORIAL

Executive Editor

Robert Sovereign-Smith

Writer

Mihir Patkar

Copy Editor

Infancia Cardozo

DESIGN

Sr. Creative Director

Jayan K Narayanan

Art Director

Anil VK

Associate Art Director

PC Anoop

Visualisers

Prasanth TR & Anil T

Senior Designers

Baiju N.V & Chander Dange

Contributing Designer

Vijay Padaya

45
PAGE

Fictional AI

Knight Rider, Bicentennial Man, Terminator, I Robot, Star Trek...need we say more?

56
PAGE

Intelligence Today

Ever heard of an aircraft that learns by watching other aircrafts? That's just one of the many skills of unmanned aerial vehicles (UAVs).

69
PAGE

Chatterbots

Read about award-winning and famous "Artificial Conversational Entities" a.k.a. AI agents with whom you can not only have stimulating conversations but whom you may also mistake for another human.

77
PAGE

IBM Watson

Meet the world's smartest artificial intelligence system in the world today – The IBM Watson.

83
PAGE

AI and the future

Technology has advanced to such unimaginable heights that you'd have to be very creative to think up the possible future of AI. Let's explore what the next wave of scientists will work on.

89
PAGE

Looking at AI under a new light

If it ain't broke, don't fix it...or wait...maybe you should.

© 9.9 Mediaworx Pvt. Ltd.

Published by 9.9 Mediaworx

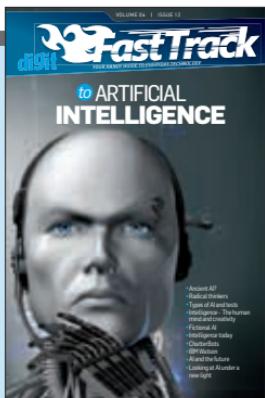
No part of this book may be reproduced, stored, or transmitted in any form or by any means without the prior written permission of the publisher.

December 2011

Free with Digit. If you have paid to buy this Fast Track from any source other than 9.9 Mediaworx Pvt. Ltd., please write to editor@thinkdigit.com with details

Custom publishing

If you want us to create a customised Fast Track for you in order to demystify technology for your community, employees or students contact editor@thinkdigit.com



COVER DESIGN: SHIGI LN

Intelligence is...

Ever since Homo sapiens learnt to communicate with each other, the one common driving force behind all you see around you is knowledge. Unlike every other species around us, mankind is hungry – too hungry, some would say. All the science and technology we take for granted today didn't exist 50 or 100 years ago, except in the realms of science fiction, perhaps. Actually, sci-fi is perhaps our greatest gift, because someone has to dream it before mankind makes it.

Mankind has shown interest in several areas of science – understanding our world, exploring space, and giving us an insight into just how minuscule and irrelevant we all are in the grand scheme of things...the one thing that still alludes us is an understanding of ourselves. Whether you're more inclined to believe Darwin or your religious books, one thing is for sure, we still can't explain ourselves as easily as Mathematics, Physics and Chemistry can explain the world we live in. Sure we can bring the recently dead back to life, or put one person's spare parts into another and save lives, but even medical science loses its exactness when it comes to the human brain. The human mind, to be more precise, or our intelligence is what makes each one of us as unique as our DNA.

Despite all the advances in medicine that you read about, such as in stem cell research, we're about as close to building a living breathing human from simple chemicals as we are to sending probes to UDFj-39546284 (the farthest known galaxy we've spotted – a mere 13.2 billion light years away). OK, perhaps we exaggerate a tad, but it's a fact that when it comes to understanding what makes us intelligent, we're still clueless... or are we?

Sure we may not be able to rebuild a working human brain from chemicals, but we may not be as ridiculously far from building a machine that can mimic one!

But what's with this verbose and rather abstract introduction, you ask? Well you're the one who picked up a Fast Track to Artificial Intelligence to

read, so you must be interested in theory and abstract concepts, or sci-fi at least.

In this Fast Track, unlike most others we do, we're not going to be helping you build anything – if we knew how to make AI, we'd be trillionaires, and you'd be buying our products. Instead, we're going to introduce the concept to you, give you a heads up about what's being worked on, and most importantly, stoke your imagination.

This book deals with everything from ancient myths about AI to actual implementations. It covers sci-fi and actual prototypes, and also introduces some varying schools of thought in the field. Most importantly, this book paints a picture of what the future might hold for all of us, or our great grandchildren, perhaps, depending on how creative or brilliant the research work is in the years to come...

Perhaps someday soon, we'll be able to finally understand what it is that makes us tick, and complete the sentence, “Intelligence is...” **d**



ANCIENT AI?

Trace the earliest references to the existence of artificial intelligence

Although the term “Artificial Intelligence” or AI is relatively new, the concept of non-human, artificially made creatures exhibiting human intelligence and emotions can hardly be called new. It can be traced back to mythological tales.

The earliest one being a love story from ancient times, purportedly written by the North African, Philostephanus about 2,300 years ago. His original story is rumoured to be the basis for Ovid's famous tale, re-written about 2,000 years ago as Pygmalion and the statue. According to the story, Pygmalion was a Cypriot sculptor who carved a statue of a beautiful woman out of ivory, and then fell in love with her. His love and desire was so strong that his pleas to the goddess Venus were answered, and she came to life.

He married her and they had a son named Paphos...intelligence and life from ivory, in this case.

It's not the only instance of artificial life through the ages. Early Judaism mythology also tells the stories of Golems, or basically creatures made from inanimate matter such as dust. Adam, of Adam and Eve fame, was said to have been made from dust, and kneaded into life. This is one of the earliest mentions of a Golem that we find in our history. The most famous of all Golems, however, isn't Adam, it's the Golem of Prague, which, as the story goes, was created by a rabbi of Prague to defend the city and people against Roman attackers. The Golem is said to have become so fierce and violent that the then-Emperor of Rome, Rudolf II, is said to have begged the Rabbi to destroy the Golem, and promised to stop the persecution of the Jews.

Now the Jews may claim that Adam, the first man, was artificial, but the ancient Greeks have a similar story about Pandora (the first woman, according to them). Pandora was made under Zeus' instructions, as revenge for Prometheus' theft of fire from him (to give to mankind). Pandora was made from earth as well, and was endowed with seductive gifts to keep men enthralled. You

may have heard about the phrase Pandora's Box and wondered what it meant. It stems from a myth about Pandora also being given the gift of curiosity, and she opened a jar unleashing all of the sicknesses and evils that have ever befallen mankind. Ancient Greece was a terrible time to be a woman, we suppose!

In more modern times, Frankenstein, by Mary Shelly is perhaps a more well-known story of an intelligence that was created by man. The most common mis-



Frankenstein's monster



No one told the bad AIs in the Terminator series about the Three Laws!

take we all make is to think that Frankenstein was the monster, whereas the book clearly states that Dr Frankenstein was in fact the creator of the “Monster”, and thus the correct usage is “Frankenstein’s monster”.

Another popular children’s tale that all of us would have read is the story of Pinocchio, the little boy made of wood, written by Carlo Collodi. Although the stories were written more as a warning for little boys to not tell lies, it’s interesting in our context that Gepetto, the carpenter, carved an intelligent, artificial creature out of wood. Of course Pinocchio yearns to be a “real boy”, and spends his life in



“Artificial” could mean anything man-made, even Pinocchio!

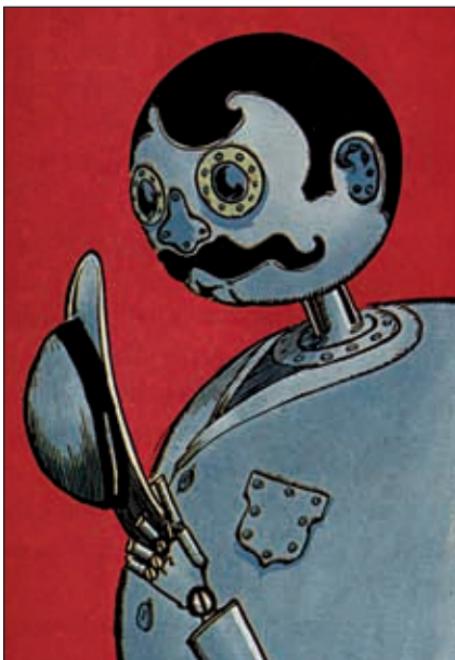
that quest, which is a common theme in sci-fi, as you'll find when you get to some of the later chapters.

The earliest reference to a mechanical intelligence, which is what all of us imagine when we think "AI" today, has to be Tik-Tok, the wind-up mechanical intelligence that L Frank Baum thought up for his Land of Oz series of books. Tik-Tok was a round copper being that had to be wound up like a watch or a clock, and had different keys that needed to be wound for thought, speech and movement. Tik-Tok is considered by most people to be the first "robot" even though that word didn't exist for a decade after his books were first published. In fact, the first known use of the word was in 1920, in a play titled Rossum's Universal Robots by Karel Čapek, a Czechoslovakian playwright, and this was a year after L Frank Baum died!

For those of you who want to point out that the Tin Man from Baum's Land of Oz series was perhaps a better fit for our imagination of AI than Tik-Tok, you should know that Baum himself said that the Tin Man was not artificially created, but in fact was a human who turned more artificial by the year as he had to have parts and organs replaced by those made of tin in order to survive.

Of course, characters from books by the likes of Asimov, Clarke, and movies such as The Terminator and Matrix have fuelled the more recent perception of AI, but those can hardly be called "Ancient", and are dealt with in a different chapter – Chapter 5, to be precise.

Going back to ancient times, we really see the initiation of AI in the process of mechanising human characteristics, or functions. One of the earliest such accounts is perhaps from Lie Yokou, a Chinese philosopher who lived around 400 BC. In a text called Liezi (meaning,

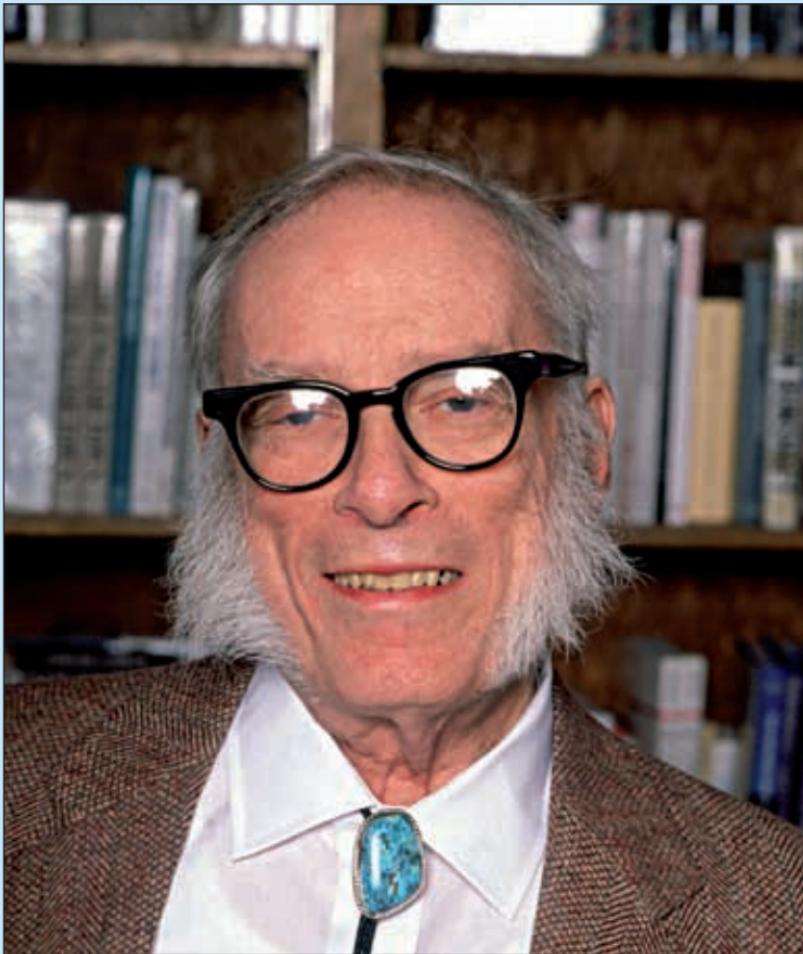


Tik Tok, from Baum's classic Oz series

“Master Lie”), which is attributed to him, he writes of an ancient (even for him) encounter of King Mu of the Zhou dynasty with a rather brilliant engineer called Yan Shi, in around 950 BC! It is said that Yan Shi brought a creation of his to show off to the king. This creation was an artificial figure that walked and sang like a human, and even winked at the ladies present. The king was fooled, and furious. He ordered the execution of Shi and his companion, until Shi showed him that his companion was nothing more than a man-made construction of leather, wood and glue, and a lot of real internal organs.

Many people also attribute the very beginnings of our quest for AI to Greek, Chinese and Indian philosophers, who first started the world on the path of deductive reasoning and logic. Many say it's the logical approach to thinking that led to the formation of the sciences and mathematics, which is the very basis for all technology – from the humble PCs of today to the intelligent beings we're trying to create tomorrow. 

CHAPTER #2



RADICAL THINKERS

The men who so strongly influenced AI
the way we know it today.

The previous chapter was dedicated to giving you just a handful of examples of why we think mankind was always headed towards advanced technology, and as a result, AI.

This chapter will look at some of the people in more recent times who were instrumental in all the progress we've seen in the field of AI. You will also be able to read about some of their work and concepts in subsequent chapters. Of course there are thousands of contributors, each of whom have played a very special part in the development of AI, and detailing them all would be beyond the scope of this book, so we'll stick to our picks of the important few.

To start talking about AI in the modern sense, one has to start with Alan Turing. Turing carries the distinction of being known as the father of computer science and AI. During the World War II, Turing was known to be the brains behind Bombe – the electro-mechanical machine that was finally able to decrypt the German Enigma code that had foiled the allies for much of WWII. However, in the context of this book, it was Turing who first posed the question “Can machines think?” in a 1950 paper titled Computing Machinery and Intelligence. Although Turing is revered by us geeks, he was ostracised in his own time for his sexual orientation. Arrested for being a homosexual in 1952, he was given a choice between prison and chemical castration. He



Alan Turing: the man behind the test we use even today to discern whether an “intelligence” has been created

chose the latter, and was forcefully given strong doses of female hormones. In 1954, he was found dead due to cyanide poisoning, and is considered to have committed suicide. In 2009, thanks to an online petition by geeks from all over the world, then British Prime Minister Gordon Brown finally issued a posthumous public apology to Turing for the way the British government had treated him over 50 years ago.

Turing is also very famous for coming up with the Turing test. At a time where there weren't any computers in existence that could even partake in such a test, Turing asked the question, "Are there imaginable digital computers which would do well in the imitation game?" Even today, the Turing test, or tests that are inspired by his original thought are what we subject AI candidates to in order to tell if they truly are intelligent or not. You can read more about the Turing test in the next chapter.

Although Turing was obviously the most important human mind behind the beginnings of artificial intelligence, it was the American John McCarthy who coined the actual term "Artificial Intelligence" that all of us use today. McCarthy also developed the Lisp (List Processing) programming language, which is the second-oldest high level programming language ever, and was pretty instrumental to AI development in the early days. Although he has many accomplishments, one in particular we want to highlight; in 1961, before many of us were born, and certainly before all of us had even had a chance to use a computer, McCarthy said in a speech that computer time sharing technologies, and even specific applications would one day be sold as a utility – much like water and electricity. Although enough people probably thought he was crazy, over 50 years later, all of us are inundated with reminders of cloud storage, cloud hosting and cloud computing in general – talk about visionaries!

Other big names in early AI research include Marvin Minsky, Allen Newell and Herbert Simon. Minsky co-founded Massachusetts Institute of Technology's Computer Science and Artificial Intelligence laboratory with John McCarthy in 1959. He was among the first to analyse artificial neural networks, and his latest book titled *The Emotion Machine* (2006) is something we would recommend to everyone interested in the subject of AI.

Herbert Simon was a multi-talented individual who contributed to cognitive psychology, AI, computer science and sociology, amongst

many others. He is one of the distinguished people to win the Turing Award. He also partnered with a much younger Allen Newell to come up with the first AI program ever – the Logic Theorist – which was able to use logic to prove or disprove a mathematical theorem that was presented to it.

Allen Newell, Simon's partner for the Logic Theorist, was a computer science researcher at RAND Corporation, and later worked with Simon at Carnegie Mellon's School of Computer Science. He was a co-winner of the Turing award (with Herbert Simon).

No such listing could ever be justified to not mention the two sci-fi greats of all time – Arthur C Clarke and Isaac Asimov.

Sir Arthur Charles Clarke lived most of his adult life in Sri Lanka, where he immigrated to from his home country – UK. He wrote many sci-fi novels, and most notably, the Space Odyssey series that introduced us to HAL, the artificially intelligent computer. He is also the person behind the phrase, “Any sufficiently advanced technology is indistinguishable from magic”, and was the first to imagine giant space elevators that would carry people from ground level up into space, and thus make rockets obsolete.

Isaac Assimov is one of the most prolific writers of all time, and is said to have written or edited over 500 books. He is most famous for coming



Isaac Asimov's Three Laws were designed for Robots, but can be applied to any AI that we create in the future



Sir Arthur C Clarke, the writer behind the famous HAL 9000

up with the concept of a “positronic brain” for AI, the word “robotics”, and also the now famous “Three Laws of Robotics” that are meant to protect human beings from being harmed by artificially intelligent robots.

You can read more about these great authors in the chapter titled Fictional AI (chapter 5). [d](#)



TYPES OF AI AND TESTS

A brief look at how artificial intelligence is broadly classified and the Turing Test used to judge an intelligent agent

As a subject taught in colleges, artificial intelligence is extremely vast, largely due to its young age. Think about it this way: what was the field of physics before Newton and Einstein had a crack at it? That's the kind of nascent stage that the study of

artificial intelligence finds itself in. The entire school of thought hasn't yet been unified, let alone have base principles, laws and theorems laid down.

Still, there are some schools of thoughts and practices that have received widespread nods of approval. While every single person working on AI might not agree with these, they are the most relevant approaches for anyone looking to understand the subject.

The two types of A.I.

Broadly speaking, there are two main branches of artificial intelligence: Strong AI and Weak AI. But the definitions of Strong AI and Weak AI are heavily dependent upon the philosophy of artificial intelligence. Are we trying to recreate an intelligent agent that can respond to a task by perceiving its environment? Or are we trying to recreate a mechanical equivalent of the human mind with all its complexities?

Let's take each of them individually to understand what they are all about.

► **Strong AI:** Every bit of science fiction you've read or seen – right from *Iron Man* to *Hitch Hiker's Guide to the Galaxy* – focuses on Strong AI. In essence, Strong AI is about building a machine that is capable of emulating the human mind to such an extent that it matches or exceeds it. As author and futurist Ray Kurzweil puts it, strong AI is machine intelligence with the full range of human intelligence.

Strong AI requires the intelligent agent to be able to think freely, or rather, outside the programming guidelines given to it. Of course, anyone reading this article will probably not be around by the time such technology comes into practical existence, but it's important to understand the goals of a science to know where it's headed and appreciate its present state.

This ability to think freely is what defines the human mind. The pinnacle of artificial intelligence would be emulating the human ability to not be bound by the predefined rules in our heads, question them and think beyond them.

► **Weak AI:** More actively in use today, Weak AI is a system that uses a set of pre-programmed rules to apply them to any task, to reach a successful fulfilment. Put simply, it's about a large and complex set of "if this, then that" rules of rationality.

While not as glamorous as Strong AI, the practicality of Weak AI is noticeable in several aspects of the world around us. And more impor-



tantly, it's being put into use rapidly. Take, for example, the ability of the new iPhone to schedule a reminder in your calendar by dictating to the Siri app. What's happening behind the scenes here? Siri is taking your voice input, smart algorithms are deconstructing that input and feeding it into their set of rules for the assigned task. So when the words "schedule" and "reminder" are heard in succession, the algorithm might ascertain that you want to perform a task concerning the calendar app on the phone, and activates a further set of rules specific to the calendar app. It will then go on to intelligently assess that task with other parts of the system, such as the GPS, so that the next time you're near a place where you can fulfil your task, your phone will tell you so. Taking a set of instructions and applying them with correlative thought by perceiving environment is Weak AI at its purest.

Weak AI is a very logical way of doing things, but that doesn't mean it's boring. In fact, the kind of applications Weak AI can have is mind-blowing.

In one of his writings, Kurzweil speaks of a future where the rapid advances in Weak AI, genetics and nanotechnology meet to create tiny super-docs. A deeper understanding of genetics is expected to turn our body into something like a logical board, where problems and solu-

tions become easier to assess. So scientists could potentially build a breed of nanobots programmed with a Weak AI that can diagnose and treat human ailments in a jiffy. A futuristic scenario of this is injecting a few nanobots into a part of your body so that they can quickly scan and diagnose the problem and treat it immediately – all without requiring any “free” thought as such.

Every conversation about artificial intelligence eventually brings up Isaac Asimov and I, Robot, so let's take a look at Strong AI and Weak AI in that context. When you talk about Asimov's “Three Laws of Robotics”, they're designed as an over-riding counter-balance for any task a robot in the future is to perform. If the task doesn't meet one of the three laws, the robot won't perform it. This is a classic example of Weak AI. But the base plot of a robot that starts thinking beyond those three laws is a prime example of Strong AI. The artificial intelligence is thinking independently – “freely” – much in the way a human does.



Alan Turing

The Turing Test

While we use the word “think” to illustrate any process of rational or irrational reasoning, it isn't a perfect definition. Artificial intelligence often gets loosely labelled as “a machine's ability to think for itself”. But what do we mean by “think” in this case?

Is it the ability to process data to come up with an output? Does that output have to meet criteria of right and wrong?



How is right and wrong determined in cases of logic? The word “think” throws up too many problems to be able to be used consistently for scientific determination, which we would all agree is necessary when trying to judge whether an artificial intelligence system is working or not. So how does one form a base test to adjudicate whether a system can be said to be an intelligent agent?

British computer scientist Alan Turing decided to redefine the parameter of what made a system an intelligent agent. Till 1950, the idea of artificial intelligence was to ask the question, “Can the machine think?” Turing came up with a new question in a 1950 paper titled Computing Machinery and Intelligence: “Are there imaginable digital computers which would do well in the imitation game?” The “imitation game” here is the ability of a machine to successfully imitate human thought and response, thus providing a much more measurable test for artificial intelligence systems.

His reasoning was that the efficacy of an artificial intelligence system lies in being able to be as human-like in its process of thinking as possible – as long as that end goal was being met, the mechanics of how it was achieved would be meaningless.

The two types of Turing Test

The “imitation game” that Turing referred to in his paper is an old party game, usually played by three players. For the sake of example, let’s call them A, B and Interrogator. In the game, A is a woman, B is a man (or vice versa) and the Interrogator – who can’t see either of them – has to guess the sex of both A and B by asking them a series of questions. It’s all done through a cold text interface like a computer screen so as not to provide any hints.

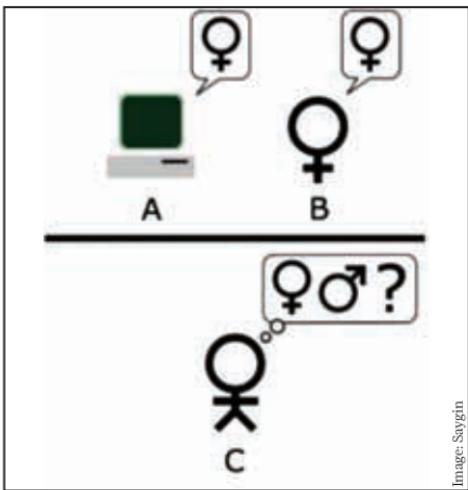


Image: Saygin

The “standard interpretation” of the Turing Test, in which player C, the interrogator, is tasked with trying to determine which player - A or B - is a computer and which is a human. The interrogator is limited to using the responses to written questions in order to make the determination.

There's another peculiar aspect of the game. Either A or B tries to mislead the Interrogator, while the other tries to provide clues to make it more obvious.

While that's the basic imitation game, Turing's wording in his paper in the 1950s left a lot to interpretation, due to which, there are now two variants of his proposed test.

► **The Imitation Game Interpretation:** The first of the two interpretations is quite akin to the original imitation game. In this, the artificial intelligence system takes on the role of A or B, while keeping the rest of the game as it is. So for example, if A is a woman and B is a machine masquerading as "male", A's objective would be to mislead the Interrogator into thinking A is a man and B is a woman; whereas B's objective would be to lead the interrogator into thinking that A is a woman and B is a man.

In this scenario, the Interrogator is never trying to determine which of the two is the artificial intelligence system. His sole purpose is to still determine the gender of A and B.

To judge the success of the AI system, the imitation game would have to be played a number of times – half where a machine is involved and half where A and B are both humans. If the human-machine duo achieved the same number and pattern of right or wrong guesses from the Interrogator as the human-human duo, then the AI would have to be deemed to be human-like, since it has gotten the same response as real people.

Later, Turing even proposed another experiment in which both, the machine and the human were trying to mislead the human to a wrong answer. But that's besides the point because the method, objective and inference of the test would be quite similar.

► **The Standard Interpretation:** Along with the Imitation Game interpretation, there is another way of looking at Turing's original paper. Called the Standard Interpretation, this is the more widely known and acknowledged version of the Turing Test. In it, the base objective is simplified to make the Interrogator figure out which of the two is a machine.

To elaborate, in this version, one out of A and B is a machine while the other is a human. Neither is trying to mislead nor give clues to the interrogator, who has to ask a series of questions to determine which of the two is human and which is the machine. If the interrogator

cannot successfully distinguish between the two, the machine would then have appropriately imitated a human enough to fool another human – thus fulfilling the test of artificial intelligence.

Pros and cons of the Turing Tests

Both types of tests have their own pros and cons, with staunch supporters and detractors in each camp.



With the Imitation Game Interpretation, the test is about the AI's ability to fulfil more tasks in terms of complexity. It is not just responding to the Interrogator, but it is also acutely aware of its own orientation as male or female and trying to either lead or mislead the Interrogator through its answers. Also, the criteria of its success doesn't lie in the Interrogator identifying it as a computer – the criteria is in being evaluated against how a human fared in a similar test. Some say that this added depth of testing provides for a more accurate judgement of whether a machine is "human-like" or not. But then there's the added caveat that the computer's task was to fool the Interrogator, not to imitate a human as closely as possible.

With the Standard Interpretation model, the test is about the AI's ability to be pitched against a human in real-time to see whether it can successfully fool the Interrogator. The AI's test here is to be as "human" as possible, and since the external influences are lesser, supporters of this type of testing claim it provides a clearer result. However, this doesn't check the AI's ability to outwit the Interrogator – it's only about providing answers without any further nuances.

But no matter which type of test the AI is subjected to, the aspects that get checked remain the same – which, after all, is the whole point of a test.

What's the purpose of the Turing Test?

As with any scientific system, the journey of the Turing Test is just as important as its conclusion. The test is so widely acknowledged as a fair

meter of judging the effectiveness of an artificial intelligence agent because it manages to test several parameters within a simple methodology.

In their book *Artificial Intelligence: A Modern Approach*, Stuart Russell and Peter Norvig write about the four base criteria that the Turing Test puts any AI through:

- ▶ **Natural language processing:** The most essential part of any artificial intelligence system would have to be its ability to converse with humans in a way that doesn't make it obvious that the other party is a machine. For this, the AI needs to both understand and reciprocate natural language. After all, what use would it be to have a robot who doesn't understand half of the things you're saying? Natural language processing is perhaps the most difficult part of building an AI agent, since it would have to take into account various dialects and slang, along with having to correlate different types of data to come up with the meaning of what was said.
- ▶ **Knowledge representation:** During a conversation, there's often a give-and-take of new information among humans. A successful AI has to have the ability to store and retrieve data it has just received from a new source, because that's key to carrying on a conversation. While it sounds difficult, this is an aspect that has advanced enough in most mechanical systems now.
- ▶ **Automated reasoning:** Reasoning, or the ability to draw rational conclusions from different data, is one of those problems that seems simple on the surface but is extremely complex the more you get into it. To store data on a system isn't a problem at all. But to give the system the ability to figure out which data correlates with which other data? That's where things start becoming complicated. And when you add the new information learned from knowledge representation to it, it's suddenly a Herculean task to come up with a system that can use old data to come up with new solutions.
- ▶ **Machine learning:** A related part of automated reasoning, machine learning is the next step in what would simply be called "pattern recognition". As humans, we have an innate ability to look at new things and make them fit into our understanding of the world, based on the patterns of information we have experienced in our lives. It's something we do without even thinking about it. But for a machine, there isn't an intuitive ability to apply patterns and detect new environments – so to pass the Turing Test would be a major triumph in that department.

The Total Turing Test

In the early 1990s, cognitive scientist Steven Harnad added that the goal of the Turing Test should be expanded for modern AI systems to include two extra components. His position is that the Turing Test takes into account a “cold” scenario where interaction between the Interrogator and the intelligent agent is through a pure text scenario. But a modern AI, especially one in a robot, would need to pass tests of sight and physical interaction to truly be deemed “intelligent”. Russell and Norvig simplify these as computer vision and robotics.

- ▶ **Computer Vision:** In the Total Turing Test, the AI would also need to pass a test of sight. In this case, while the Interrogator can't see A or B, he would be able to test their vision in some way to ascertain whether they're human or machine. For an eventual intelligent robot, this would need to be a crucial element – and it would go on to test the knowledge representation aspect of the system a lot more thoroughly.
- ▶ **Robotics:** While machine learning is tested in the form of detecting new environments, the Turing Test doesn't check the AI's ability to manipulate these environments. How can a modern AI system – used so frequently to operate heavy machinery – be considered complete without the ability to interact with its environment in a way that most humans would?

Not An Exact Science

While it's the most widely accepted form of testing artificial intelligence, the Turing Test can't be called an exact science. It doesn't give any denomination to intelligence nor any measurable capacity to pit one AI against another AI. To put it bluntly, the Turing Test is “all or nothing” in the sense that it can judge a machine to be artificially intelligent or not.

This obviously begs the question of what is intelligence. If a machine is more intelligent than any other animal on the planet, but less intelligent than a human being, is it still intelligent? And if a machine does pass the Turing Test and several others do the same later, then which is the “smartest” machine of them? The lack of a quantitative system in the Turing Test makes it an inexact scientific tool.

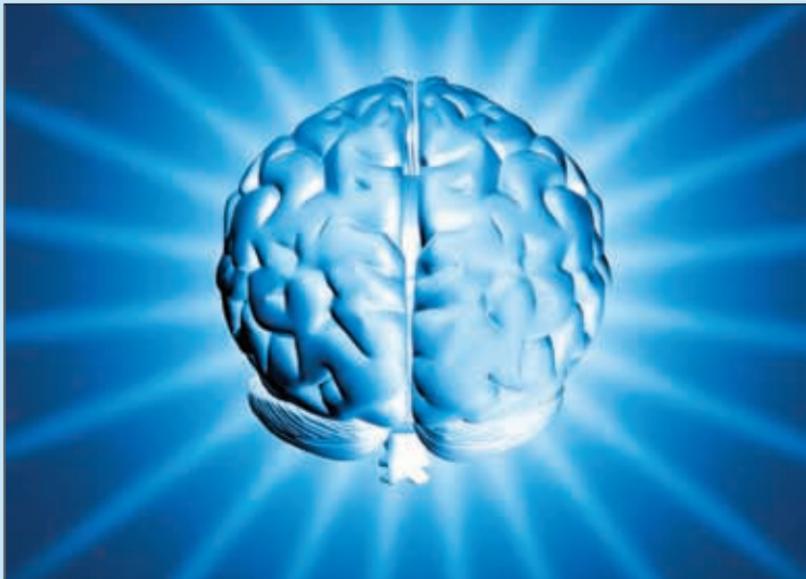
Still, currently, there aren't any alternatives either. The closest anyone has come has been a team of scientists who are trying to come up with a “universal intelligence system”.

Universal Intelligence System

Australian computer scientist David Dowe of Monash University and Spanish researcher José Hernández-Orallo of the Polytechnic University of Valencia came together to devise a new system to measure intelligence itself. Their work is based on a theorem called the “Kolmogorov complexity”, which basically states that any object’s complexity is a measure of the computational resources needed to express it. And the minimal resources needed in doing so is the Kolmogorov complexity of that object.

To put it in layman’s terms, the amount of data needed to describe an object is how complex it is. Of course, each object has different ways by which you can describe it. A “red book” can also be called a “collection of papers bound together in a casing whose appearance is closest to the longest wavelength of light as can be seen by a human eye”. In this example, the object’s Kolmogorov complexity would be “red book”.

Of course, for testing intelligence, you would use different objects like strings of characters or binary. And that’s what Dowe and Hernández-Orallo are doing. By having a quantifiable measure for intelligence in the form of the Kolmogorov complexity, they’re able to judge whether one AI system is more intelligent than another. 



INTELLIGENCE – THE HUMAN MIND AND CREATIVITY

How the different aspects of the
human brain – reasoning, knowledge
and creativity – affect the making of
artificial intelligence

The human race has always been proud of its brain. And rightly so. Hey, if we were smart enough to invent chocolate, we obviously have something going for us, right? Call it hubris or just plain rationality, but the success of artificial intelligence

systems has hence come to rely on the human brain as a yardstick. As we saw with the Turing Test, an intelligent agent is deemed so only if it manages to fool a human into thinking it's human.

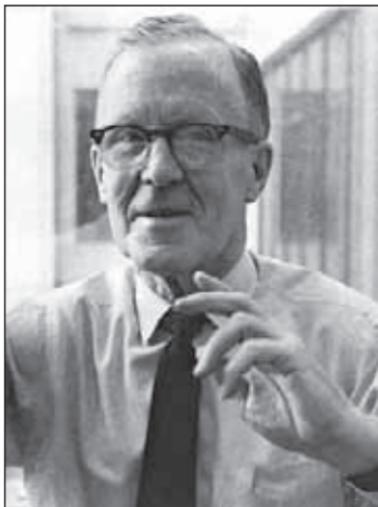
With that goal always at the back of their minds, scientists have sought to emulate the human mind in machines, albeit with different paths and results. The earliest form of these efforts meant studying computational neuroscience – the science of how the human brain and nervous system processes information – and applying it to electronics.

A lot of the initial work was based on the findings of Donald Hebb, a neuropsychologist at the McGill University in Canada. In the 1930s, he was one of the first to talk about “learning” from a neural perspective. Hebb said learning was all about the neural networks i.e. the clusters of cells and neurons. The strength of the connections between these cells and neurons affected how a human brain learns. To this end, he theorised a mechanism by which the brain learns to distinguish objects and signals, add and understand grammar.

Hebb's followers took his ideas and reshaped them into two broad concepts:

- i) Autonomous systems:** Each neuron is a self-adjusting cell that changes the strengths of its connections to other neurons when learning. This way, fewer errors are made when it repeats a task.
- ii) Instantaneous learning:** Scientists also presumed that learning in the brain is “instantaneous”, i.e. as soon as something to be learned is presented, the appropriate brain cells use their “local learning laws” to make instant adjustments to the strength of their connections to other neurons. When learning is complete, the brain discards its memory of the learning example.

Neurology became the focus of researchers in the artificial intelligence fields and soon, they started building machines based on the wiring of



Donald Hebb

the human brain. Of course, it wasn't like they were creating a shape like a brain and putting wires into it. Think of it more like a blueprint that formed the basis of the structure of early artificial intelligence devices. However, these quickly failed.

Dr. Asim Roy, a professor of Operations Research and Information Systems at Arizona State University, points out that these models are fundamentally flawed since they don't reproduce the external characteristics of the human brain, principally its independent way of learning. Therefore, methods based on these classical ideas require constant intervention by engineers and computer scientists – providing network designs, setting appropriate parameters correctly, and so on – to make them work. This drawback is severe, he maintains.

It was a thought that struck some other scientists early in the 1950s, and the attention of AI researchers started turning towards the psychological workings of the human brain rather than the physical workings.

Simulating how the brain thinks

It's in the importance of symbols that the working of the human brain and the history of artificial intelligence start intertwining. And the origin of it all couldn't be more innocuous. Some time in the 1950s, at the RAND corporation – a global non-profit think tank – a digital map was being printed. Those weren't the days of laserjets, so this humble printer was manipulating alphabets, numbers and punctuation marks to best simulate the map. Herbert Simon was witness to this ordinary act, which sparked an extraordinary thought process in his head.

While everyone around him saw a printer doing its job, Simon saw a machine using symbols – which is essentially what those alphabets, numbers and punctuations were – to convey a meaningful message in a medium it wasn't intended for. He was quick to realise that a machine that could manipulate symbols could potentially simulate the human thought process.

This realisation had a lot to do with Simon's own background as



Herbert Simon

well. As a political scientist, Simon was deeply involved in the understanding of how large decision-making systems worked. In his research, he had realised that making a decision wasn't a mathematical process. Mathematicians took the human brain's decision-making ability and confused rationality with optimisation. The "right" decision was always the outcome that would be most fruitful given a situation's parameters. Simon, however, proposed that rationality wasn't always about getting the optimal output. Instead, he said that it involved a wide variety of factors, including the limitations of a person's mind, the amount of time they have to make a decision and the information they have at that point. He contended that it was impossible to have perfect and complete information at any given time to make a decision – a perspective that would eventually win him a Nobel Prize. Simon would go on to encapsulate all of this into a theory called "Bounded Rationality", but going into that would make us digress too much from our core matter here.

What concerns here is Simon's view that decision-making need not always be about coming up with the most optimal solution each time. It was this leniency afforded by Simon to the human thought process that allowed him to see the printer in a way that no one around him did.

In Allan Newell, he met a kindred spirit. Newell was a computer scientist at RAND who had developed the program to print the map.



John McCarthy set the ball rolling with a simple question

Studying organisation theory, Newell's interest was similar to Simon's and together, they started chalking out a way to build a machine that could "think" in the way they defined the term "thinking". It wasn't about the neural framework of the human brain, nor was it about coming to the best possible solution in any situation. It was based on how the human brain thought – make the best decision given the information at hand and the limitations enforced.

Still, they needed to formalise this into a science. The duo set about creating programs and carrying out studies that sought to understand how humans figured things out, i.e. the psychology involved in problem-solving. This set about a revolution in cognitive science, which sparked an interest in a branch of artificial intelligence called "Cognitive Simulation" – getting machines to mimick the way the human brain works in solving problems.

Simon and Newell had a very systematic approach to this newfound branch, carrying out meticulous psychological studies and experiments to come up with patterns and techniques people used to solve problems. These were then digitised into computer programs to help artificial intelligence systems "think" for themselves.

Human logic and applying it to AI

While the world of AI was firmly changing its focus to emulating the inner workings of the human brain, the psychology of the mind is so vast that there were several different schools of thought – no pun intended. Newell and Simon preferred to focus on techniques of problem-solving, but John McCarthy, a cognitive and computer scientist at Stanford University in the US, decided to look at it a little differently and asked AI a question we all seem to have asked someone in our lives: "Don't you have any common sense?"

Common sense might be a subjective term, but the basis of what McCarthy was talking about was logic. The human mind's capacity for logic is unparalleled in nature. And yet, it wasn't formalised as an ability till Aristotle came along and started compartmentalising and defining aspects of the human psyche.

Logic can be something as simple as a syllogism (All men are mortals. All Greeks are men. So All Greeks are mortal.) or something a bit more complex like deductive reasoning. It's when it comes to reasoning that McCarthy's contributions are highlighted.

Simon and Newell were busy figuring out the technique with which humans solved problems, and coming up with ways to turn that into digital

programs. McCarthy decided that this was just the tip of the solution – the deeper necessity lay in finding out the essence of reasoning and logic.

A lecturer in MIT tried to explain it with the use of a riddle. Let's try it out. In a large maze, there are two different paths to reach the finish line on the other side. Each path poses a different hurdle – one is longer and safer, the other is shorter and more treacherous. Which of the two paths would you choose and why? Got an answer in your head?

The students in the class were asked one by one what they chose and their reason for doing so. Invariably, there were several different thought patterns emerging. And that's when the professor sums up the lesson: No matter which option you go for, you'll end up reaching the same finish line. It doesn't matter when you reach it or how you reached it, since that wasn't the part of the problem.

Similarly, logic works in a way where each problem can have several different ways of solving it. McCarthy wanted to study the essence of abstract reasoning and problem-solving to emulate this ability. For there to exist artificial intelligence, he contended, the machine needed the ability to think for itself while basing the foundation of the thought process on the core of the human reasoning process. To put it differently, he didn't want to emulate a certain “human algorithm of reasoning” in digital format, he wanted to emulate the ability of human reasoning itself.

To this end, McCarthy dove into the world of logic and came up with several pioneering theories and applications, in the process solving some monumental problems in the world of knowledge representation. Given his ability to break down the core concepts of human psychology, it's little wonder then that this is the man who came up with the term “artificial intelligence” itself.

Knowledge and intelligence

McCarthy's reliance on logic and Simon's reliance on the technique to solve problems both had one common hindrance: knowledge. At the time they were conducting their experiments, the storage capacities of computers severely limited their options. Indeed, their entire line of thinking could be said to be a workaround for the limitations of storage space.

But in the human brain, storage space is one of the most important aspects of how we function. Knowledge is power, goes the old adage, after all. The human brain stores massive amounts of information in a physically small space, thus making it an exponentially useful tool. This

knowledge base can be called upon quickly and efficiently as per the requirement of any task – be it problem-solving, reasoning or anything else. Indeed, even the ability to solve one problem and use its logical inference to tackle another problem requires “knowledge” in the form of stored information about a certain way of doing one thing.

In the 1970s, there was a turnaround in the kinds of storage spaces available to computer scientists. The human brain tackles problems by using two aspects: knowledge and reasoning. As large amounts of memory started becoming available, researchers started dividing the artificial intelligence systems into two such halves as well.

Stanford University’s Edward Feigenbaum was at the forefront of this movement, developing what is the world’s first “expert system”. An expert system is essentially an intelligent agent that relies on two parts: the static inference engine and the dynamic knowledge database.

The inference engine is basically what Simon, Newell and McCarthy worked on – the ability of a computer to reason. For an expert system, it didn’t matter what kind of logical thinking it was based on, as long as the computer was able to reason. The real genius of this lied in its ability to use a knowledge database to store information as well as add to it over time. This way, the expert system was using reasoning capabilities with its pre-existing knowledge, making it more like a human brain than ever before. And over time, boffins developed the ability to add to the dynamic knowledge base, making it such that the AI could draw on previous logical inferences to further expand its knowledge base and thus increase the speed of reasoning.

Neat vs Scruffy

While the theories of Newell, Simon and McCarthy were being taken as gospel in most scientific communities, a new line of thinking came up in the mid-70s. Roger Schank, a cognitive scientist at Yale University, had a different interpretation of what was needed to solve the massive problem of coming up with an emulation of the human brain. His solution? Nothing.

Well, that’s putting it a bit harshly, but Schank’s premise was that the human mind is too vast and too complex to be reduced to a system that can be applied in any scenario homogenously. In his studies, he had found that different tasks required different parts of the human brain to kick in; and more importantly, people had different parts of the brain kicking in for the same task at different times or situations. So then how

could such a complex mechanism be reduced to a “neat” solution like the one proposed by McCarthy or Simon? The idea of boiling down “logic” into a semantic network didn’t fit in with the idea of the human brain being unimaginably complex.

Instead, Schank found himself agreeing with the approach employed by cognitive scientist Marvin Minsky of the Massachusetts Institute of Technology. Minsky and his team were the original “hackers” – it was in their lab that the term was born. Their idea was that each task required a specific solution and there’s no “one ring to rule them all” universal principle for artificial intelligence. Minsky’s team used to work on each problem individually, tweaking code – or “hacking away” – till they got it to work right. Schank described this philosophy towards AI as “scruffy”, in stark contrast to the “neat” solutions.

The result of the scruffy approach was that it made more rapid strides in the non-problem-solving aspects of AI than neat solutions. For example, MIT computer scientist Joseph Weizenbaum created a computer program called ELIZA, which has often been described as the first truly functional natural language processing automation – or rather, “chatbot”. Acting as a psychotherapist, it gave remarkably human-like answers to questions. For instance, if the user typed “My mother hates me”, ELIZA would respond with “Who else hates you?”

A lot of modern robotics use scruffy AI rather than neat AI because of the way the robots are meant to behave. A long sequence of logical formalisations would lead to more resource-intensive behaviour, which would in turn be slower. If the robots were instead programmed to perform specific tasks based on reacting to a certain stimulus – which is essentially the scruffy approach – they would function in a more practical manner.

Embodied and Computational Intelligence

Every single school of thought regarding AI was obsessed with the human brain and eventually, they all started hitting dead ends. In 1986, Rodney Brooks – a professor of robotics at MIT – wrote a monumental paper called *Elephants Don’t Play Chess*, in which he stated that reasoning about the physical world is much simpler than interacting with it. He started focusing his work in AI on creating biologically-inspired robots.

While this is seen as the point where computational intelligence took off, its roots can be traced back a few more years to the emergence



Thinking outside the box

of embodied cognition – a field of science that believes that human intelligence can't be boiled down to the functioning of the brain, and that higher intelligence is impossible without perception, movement and visualisation.

This approach quickly led to the emergence of computational intelligence, which is a massively popular branch of AI today, which is essentially an amalgamation of the best aspects of pre-existing techniques. It borrows from “neat” theories, “scruffy” theories, fuzzy logic systems and even neural networks.

The most important breakthrough in computational intelligence came in 1986, when American psychologists David Rumelhart and James McClelland wrote a paper called *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*. In it, the duo described in great detail the inner workings of the human brain in association with perception and created computer simulations of perception. This was the first time that scientists could test neural processing, thus making AI more true-to-life.

Essentially, computational intelligence took the field of AI outside of “the human brain” and made it about “the human mind” in all its

different aspects – a field so large that you'd need to be an expert before getting further into it.

Artificial Intelligence and Creativity

The human mind is a complex mechanism – a kind of super computer that manages to combine knowledge and reasoning to tackle gargantuan problems in a jiffy. And yet, understanding these aspects of the brain has helped scientists come to a point where they feel confident of claiming that at some point, a computer will be able to carry out these functions. But intelligence has one more factor whose inclusion throws the world of artificial intelligence into a tizzy: Creativity.

What is creativity?

Creativity would have been a tricky thing to explain if we weren't all born with it. The ability to think of something new, or applying old formulae in a new way, separates the mind of the human from any other animal.

Of course, this is a simplistic explanation of an extremely complex matter. Think about it this way: if we didn't have the ability to be creative, would we intrinsically know what is creative and what isn't? Take writing, for instance. The ability to write fiction is often considered a creative output. The same can be said about art – a field that is most often associated with creativity. However, if an artist is replicating a scene he saw somewhere, then is it still creative? And if a writer is creating a new way of expressing an old idea, is it still creative? It's incredibly difficult to define creativity in a universal manner because the term holds different meanings for each human.

For the sake of this write-up, however, let's treat creativity the way one of the foremost thinkers on the subject, Margaret Boden, defines it: "Creativity is the ability to come up with ideas or artefacts that are new, surprising, and valuable. "Ideas", here, includes concepts, poems, musical compositions, scientific theories, cooking recipes, choreography, jokes ... and so on, and on. "Artefacts" include paintings, sculpture, steam-engines, vacuum cleaners, pottery, origami, penny-whistles ... and you can name many more."

Why is creativity important in thinking?

Taking the very base meaning of being original, creativity is still a large and complex issue. But its need in the thought process cannot be understated.

Take some of the most important thinkers of our times. Sir Isaac Newton, as we all know, discovered gravity when an apple dropped on his head. When put so simply with our modern understanding of gravity, it seems like an everyday occurrence and only natural that he thought of gravity from that point. But try and put yourself in his shoes, in a time when the concept of gravity did not exist. Without using the creative muscles of his brain, Newton could never have thought outside the box of then-prevalent scientific reasoning to come up with a new theory about what made the apple fall from the tree.

For a modern example, let's look at the late Steve Jobs. Love him or hate him, you can't deny the fact that he spearheaded several projects that redefined technology, like the iPod. And let's face it, it wasn't that Jobs revolutionised technology by inventing some new metal or discovered a new process. He used his creativity to solve the problem of using existing technology in the field of portable music. It was a new way of thinking about an old problem – and such creativity is crucial to problem-solving.

So any modern artificial intelligence system could not be said to be true AI without the ability to think creatively. It would need to come up with solutions that have eluded the thought process of its programmers – and perhaps even escaped their subconscious.

The problem with “programming” creativity

Now that we have some sort of a loose definition for creativity, it's important to understand the basic problem that has dogged everyone in the field of artificial intelligence: if we accept that creativity is the ability to come up with a new way of thinking, how do we program it into a machine? The very nature of programming goes against this.

Programming, in essence, is laying down a set of instructions to give to a machine, which it duly carries out. Now, these sets of instructions are laid down by a human. So conceptually, the human has thought of these instructions. But how does a human instruct a machine to think of something he has never thought of? You can program the machine to think, of course, but how do you program a machine to think about something that is outside the realm of thought process itself?

More importantly, programs are task-oriented. An intelligent agent is fed a series of instructions that make it calculate and reason with a very specific objective in its digital mind. With creative thinking, what is the objective?

Computational creativity

Given the difficulty of defining creativity and employing any rough definitions into the field of computers, artificial intelligence experts sought to come up with a new field of science to further develop this conundrum. It was called computational creativity and at the forefront were the same three gentlemen who pioneered the usage of reasoning and the 'neat' approach to AI: Robert Simon, Alan Newell and J. C. Shaw.

For the sake of computing and artificial intelligence, they decided to focus on the two aspects of creativity that would be most important in an intelligent agent: novelty and usefulness. This disregards the field of creativity for the sake of art or for no real objective at all, and we'll come to that later.

When an AI is posed a problem, Simon, Newell and Shaw came up with four guidelines to define computational creativity:

- i) The answer is novel and useful (either for the individual or for society)
- ii) The answer demands that we reject ideas we had previously accepted
- iii) The answer results from intense motivation and persistence
- iv) The answer comes from clarifying a problem that was originally vague

How does this help the AI? Well, much like Asimov's three laws of robotics, these guidelines make sure that the AI will look for creativity. The first guideline (novel and useful) clarifies the objective of the solution being sought. The second guideline (rejecting previously successful ideas) insists that the AI abandon any thought process that it has already gone through to come to a successful conclusion, thus ensuring that the solution will be new. The third guideline instructs the AI to keep looking for a solution even if all earlier paths have been tried and failed, since the machine would otherwise just report a negative result. And the fourth guideline, perhaps the most intriguing, tells the AI that the problem as stated may not have completely meant as stated – i.e. make the machine re-examine the true nature of the problem and not just its semantic meaning, thus opening its horizons for different interpretations and problems.

Mimicking human creativity for problem solving

Given the parameters that Simon, Newell and Shaw laid down, research in applying creativity to computers began in earnest. There have been several programs over the years which have successfully applied this to problem-solving, but none has been as interesting as the recent work at Rensselaer Polytechnic Institute in the US.

Take this puzzle that was presented: A dealer in antique coins gets an offer to buy a beautiful bronze coin. The coin has an emperor's head on one side and the date "544 B.C." stamped on the other. The dealer examines the coin, but instead of buying it, he calls the police. Why?

Solving this "insight problem" requires creativity, a skill at which humans excel (the coin is a fake – "B.C." and Arabic numerals didn't exist at the time) and computers do not. But a new explanation of how humans solve problems creatively – including the mathematical formulations for facilitating the incorporation of the theory in artificial intelligence programs – provides a roadmap to building systems that perform like humans at the task.

Ron Sun, Rensselaer Polytechnic Institute professor of cognitive science, said his new "Explicit-Implicit Interaction" theory could be used for future artificial intelligence.

"As a psychological theory, this theory pushes forward the field of research on creative problem solving and offers an explanation of the human mind and how we solve problems creatively," Sun said. "But this model can also be used as the basis for creating future artificial intelligence programs that are good at solving problems creatively."

The paper, titled *Incubation, Insight, and Creative Problem Solving: A Unified Theory and a Connectionist Model*, by Sun and Sébastien Hèlie of University of California, Santa Barbara, appeared in the July edition of *Psychological Review*. Discussion of the theory is accompanied by mathematical specifications for the "CLARION" cognitive architecture – a computer program developed by Sun's research group to act like a cognitive system – as well as successful computer simulations of the theory.

In the paper, Sun and Hèlie compared the performance of the CLARION model using Explicit-Implicit Interaction theory with results from previous human trials – including tests involving the coin question – and found results to be nearly identical in several aspects of problem solving.

In the tests involving the coin question, human subjects were given a chance to respond after being interrupted either to discuss their thought process or to work on an unrelated task. In that experiment, 35.6 per cent of participants answered correctly after discussing their thinking, while 45.8 per cent of participants answered correctly after working on another task.

In 5,000 runs of the CLARION program set for similar interruptions, CLARION answered correctly 35.3 per cent of the time in the first instance, and 45.3 per cent of the time in the second instance.

"The simulation data matches the human data very well," said Sun.

Explicit-Implicit Interaction theory is the most recent advance on a well-regarded outline of creative problem solving known as "Stage



A collage of some of the works drawn by AARON

Decomposition", developed by Graham Wallas in his seminal 1926 book *The Art of Thought*. According to stage decomposition, humans go through four stages – preparation, incubation, insight (illumination), and verification – in solving problems creatively.

Building on Wallas' work, several disparate theories have since been advanced to explain the specific processes used by the human mind during the stages of incubation and insight. Competing theories propose that incubation – a period away from deliberative work – is a time of recovery from fatigue of deliberative work, an opportunity for the mind to work unconsciously on the problem, a time during which

the mind discards false assumptions, or a time in which solutions to similar problems are retrieved from memory, among other ideas.

Each theory can be represented mathematically in artificial intelligence models – indeed, they bear a resemblance to the computational creativity guidelines set down by Simon, Newell and Shaw. However, most models choose between theories rather than seeking to incorporate multiple theories and therefore are fragmentary at best.

Sun and Hélie's Explicit-Implicit Interaction (EII) theory integrates several of the competing theories into a larger equation.

“EII unifies a lot of fragmentary pre-existing theories,” Sun said. “These pre-existing theories only account for some aspects of creative problem solving, but not in a unified way. EII unifies those fragments and provides a more coherent, more complete theory.”

The basic principles of EII propose the coexistence of two different types of knowledge and processing: explicit and implicit. Explicit knowledge is easier to access and verbalise, can be rendered symbolically, and requires more attention to process. Implicit knowledge is relatively inaccessible, harder to verbalise, is more vague and requires less attention to process.

In solving a problem, explicit knowledge could be the knowledge used in reasoning, deliberately thinking through different options, while implicit knowledge is the intuition that gives rise to a solution suddenly. Both types of knowledge are involved simultaneously to solve a problem and reinforce each other in the process. By including this principle in each step, Sun was able to achieve a successful system.

“This tells us how creative problem solving may emerge from the interaction of explicit and implicit cognitive processes; why both types of processes are necessary for creative problem solving, as well as in many other psychological domains and functionalities,” said Sun.

Creativity and art in Artificial Intelligence

The idea of creativity is usually linked with another intrinsically human quality: Imagination. Indeed, when it comes to art, the two seem almost intertwined. And there can be no question about art's usefulness. It's no longer an arbitrary and meaningless pursuit, as several psychologists and philosophers have spoken about its therapeutic and inspirational qualities.

But does it fit into the concept of artificial intelligence? Well, yes and no. It depends on how you want to look at the field of AI. According

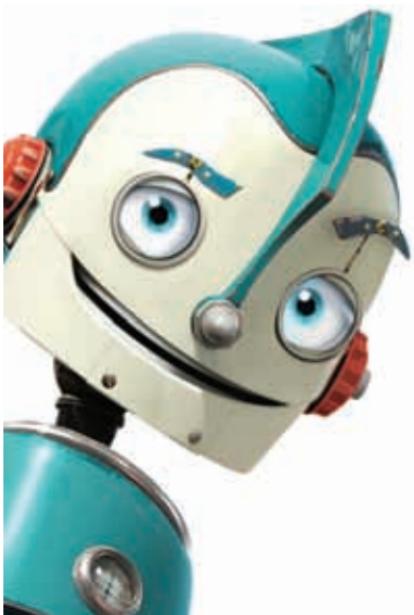
to the standard Turing Test, imagination in a visual sense would not be necessary of an AI since the entire test is conducted from a textual perspective. Even if you look at the Total Turing Test, it still needs the AI to only identify the visual cues it is given – its ability to imagine or “be creative” is never tested.

Yet, can we really have true AI without the ability to be creative and imaginative? One of the leading thinkers in artificial intelligence, Ray Kurzweil, certainly didn't think so and neither did Harold Cohen, the art professor at the University of California at San Diego. In the early 70s, Cohen created a new piece of AI that he called AARON, which was born out of the desire to answer a simple question: what are the minimum conditions under which a set of marks function as an image?

In an article for Stanford University, Cohen wrote: “On the simplest level it was not hard to propose a plausible answer: it required the spectator's belief that the marks had resulted from a purposeful human, or human-like, act – actions performed by a computer program possessing cognitive capabilities quite like the ones we use ourselves to make and to understand images. The earliest versions of AARON could do very little more than distinguish between figure and ground, closed forms

and open forms, and perform various simple manipulations on those structures. That might not have been enough, however, had AARON not performed, as humans do, in feedback mode. All its decisions about how to proceed with a drawing, from the lowest level of constructing a single line to higher-level issues of composition, were made by considering what it wanted to do in relation to what it had done already.”

What's remarkable about the program is that AARON is completely autonomous. Its intelligence lies not only in its ability to come up with art, but



A robot with emotion? Possible?

in the fact that there's no human input whatsoever at any stage of the "creative" process. Once AARON starts drawing, a human can't influence it in any way. In fact, the only thing that can influence it is AARON itself.

Of course, there has to be human input at some point for AARON to "see" the outside world. Over time, Cohen added several models along a three-dimensional axis to AARON, thus making it more aware of the visual world. In fact, he stated that in one instance when Cohen was developing the knowledge of structure of heads and faces to make AARON draw a facsimile of a friend, the AI started generating make-believe people that looked eerily like some of the people Cohen knew.

AARON has the unique ability to work on feedback. What this means is that each single line that AARON draws is registered by the machine, and it goes through a series of thousands of decisions to decide what to draw next. Of course, given modern mechanics, this happens quite a lot faster than it sounds.

But there's still a fundamental difference between the way AARON draws and the way a human draws. For AARON, any painting starts with the foreground object and works its way backwards to the background. Given the feedback system it works on, the "whole" picture isn't seen from the start like a human sees it. Indeed, human artists have a "vision" that they try to replicate on canvas, whereas AARON uses a systematic, step-by-step approach to reaching a full picture at the end, even though it might not know what that picture is going to be at any point of the creation process.

Kurzweil and Cohen have also worked on other programs about the intersection of creativity, art and computers, notably a robot that generates poetry. Similarly, composer and computer scientist David Cope's EMMY could write beautiful symphonies reminiscent of the works of Bach and Mozart.

Can AI have a sense of humour?

Getting back to the core conversation about creativity, it's interesting to know that its usage in the human brain is not just in creating something new, but also in understanding the link between two seemingly disjointed artefacts. A sense of humour is perhaps the greatest example of this.

The human ability to have a sense of humour is dependent on understanding how one object is linked with another. As Margaret Bolden explains:

“Combining ideas creatively isn’t like shaking marbles in a bag. The marbles have to come together because there is some intelligible, though previously unnoticed, link between them which we value because it is interesting – illuminating, thought-provoking, humorous ... – in some way. We saw also that combinational creativity typically requires a very rich store of knowledge, of many different kinds, and the ability to form links of many different types.”

“And we don’t only form links, we evaluate them. For instance, we can recognize that a joke is “in bad taste”. In other words: yes, the links that the joker is suggesting are actually there (so it’s a real joke). But there are other links there also, which connect the ideas with sorrow, humiliation or tragedy. The joker should have noticed them, and should have refrained from reminding us of them.

“For a computer to make a subtle combinational joke, never mind to assess its tastefulness, would require (1) a data-base with a richness comparable to ours, and (2) methods of link-making (and link-evaluating) comparable in subtlety with ours.”

The one artificial intelligence system that continually comes up in discussions about joke-making is JAPE (Joke Analysis and Production Engine). Developed by British computer scientists Graeme Ritchie and Kim Bininstead of the Edinburgh University, JAPE has the ability to tell the kinds of jokes you would find in an old *1001 jokes to tell your school friends* book.

In scientific terms, it generates punning riddles, from a linguistic model of puns, including a collection of high-level pun schemata, and a non-humour-specific lexicon w/ semantic, syntactic, phonological, and psycholinguistic information. JAPE is essentially an intelligent agent that’s contextually aware and has natural language processing capability to produce puns.

Here are a few examples of the kinds of puns JAPE generates:

“What kind of E.T. drinks beer?” “An ale-ian.”

“What do you call a breakfast food murderer? A cereal killer.”

“What’s the difference between a dog and a painter? One can shed a coat and the other can coat a shed.”

“What’s the difference between a pretty glove and a silent cat? One is a cute mitten and the other is a mute kitten.”

In a field test, the researchers made a collection of JAPE jokes, human-made jokes and non-jokes of a similar part. These jokes were shown to

a group of 100 children who were asked to evaluate what was a joke and what wasn't.

The kids classified 60 per cent of the JAPE jokes as jokes, 80 per cent of the human-made jokes as jokes, and 20 per cent of non-jokes as jokes. After adjusting for the vocabulary of the children at their age, there was no significant difference between the number of jokes classified as "jokes" for JAPE-made and human-made.

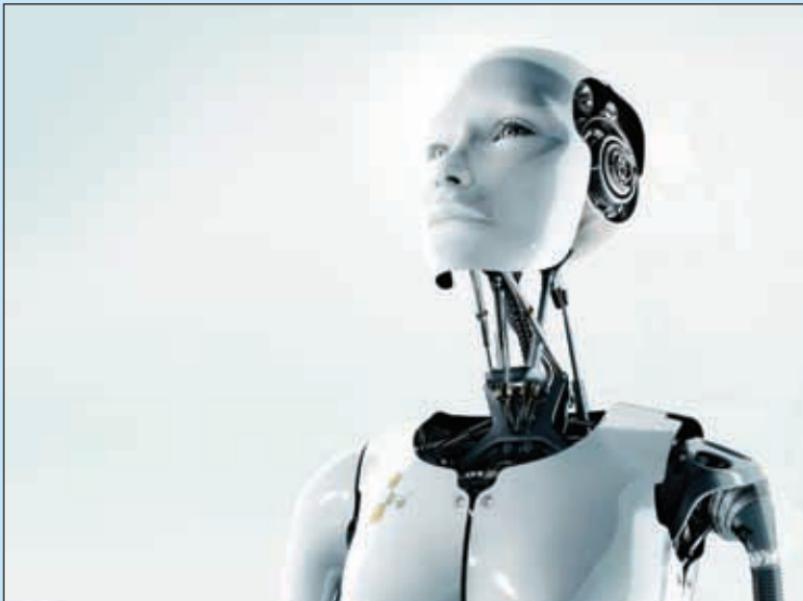
While JAPE was a successful early prototype of proving that machines could eventually develop a sense of humour, they still have a long way to go before artificial intelligence becomes aware enough to figure out the parameters for context on its own.

Why the human mind shouldn't be a model for AI

We've gone into much detail about how artificial intelligence systems use the human mind as a blueprint for creation and measurement. But how relevant is it when you actually think about it? Should the test of an artificial intelligence system really be its ability to mimic the human mind so closely that it fools another human mind, as is required by the Turing Test? There are, of course, several arguments on each side of the issue.

Consider that classic example given by Russell and Norvig where they say, "The quest for 'artificial flight' succeeded when the Wright brothers and others stopped imitating birds and learned about aerodynamics. Aeronautical engineering texts don't define the goal of their field as making 'machines that fly so exactly like pigeons that they can fool even other pigeons'."

Some food for thought. 



FICTIONAL AI

The most important visions of artificial intelligence in science fiction and some of the science behind them

The coming of age of computers and the possibility of artificial intelligence was the kind of stuff every writer dreams of. Here was unexplored territory grounded in reality. It was in 1956 that John McCarthy first coined the term “artificial intelligence” and in the next few years, a slew of science-fiction writers took it to its extreme possibilities.

Any discussion about artificial intelligence in science fiction usually ends up in a debate between who was the better writer: Isaac Asimov or Arthur C. Clarke. Both authors have written innumerable pieces on AI that far surpass any of their peers – those born then and those still to

come. In fact, a lot of their predictions about the future have turned out to be eerily accurate. Some others have obviously failed as technology progressed, but the base ideas of their works are what capture the imagination of any reader.

Not many know that Asimov and Clarke shared a good friendship. The two often met to discuss ideas and the future, unbothered by the world's need to crown one of them as better than the other. Well actually, that's not entirely true. They were quite aware of this and even came up with a small joke in the form of the Asimov-Clarke Treaty of Park Avenue, which they compiled while on a cab ride in Park Avenue, New York. The treaty stated that when asked, Asimov would always insist that Clarke was the best science fiction writer in the world (reserving second best for himself), while Clarke would always maintain that Isaac Asimov was the best science writer in the world (reserving second best for himself). In fact, Clarke dedicated his book *Report on Planet Three to Asimov*, writing, "In accordance with the terms of the Clarke-Asimov treaty, the second-best science writer dedicates this book to the second-best science-fiction writer."

So why was their work so revolutionary and widely appreciated? Even examining a few of their ideas, one comes away with the answer.

Asimov's Positronic Brain

In Isaac Asimov's world of artificially intelligent robots, perhaps no piece of technology has been as important as the "positronic brain". Indeed, so popular was the term that it has since been replicated and reused in several different science fiction franchises. Almost every android in Asimov's robot series of novels is equipped with this computer in their heads.

But the funniest part about the positronic brain is that the technology behind it has never been truly explained. Asimov let through information in little bits and pieces in several different books and stories, but never sat down and explained the machinations of the positronic brain.

In fact, he has publicly stated that the word "positronic" was merely a way of catching people's fancy. A positron is basically the antiparticle or antimatter of the electron. To put it simply, whereas an electron has a charge of -1, a positron has a charge of +1. There's no real science in Asimov's suggestion that such a piece of sophisticated computational machinery would be built using positrons. The only reason he used the term is because the positron was a newly discovered particle when writing the robot stories and he thought the adjective form (positronic, inspired from electronic)

would conjure up a vision of a futuristic endeavour. A little disappointing when it's explained, right?

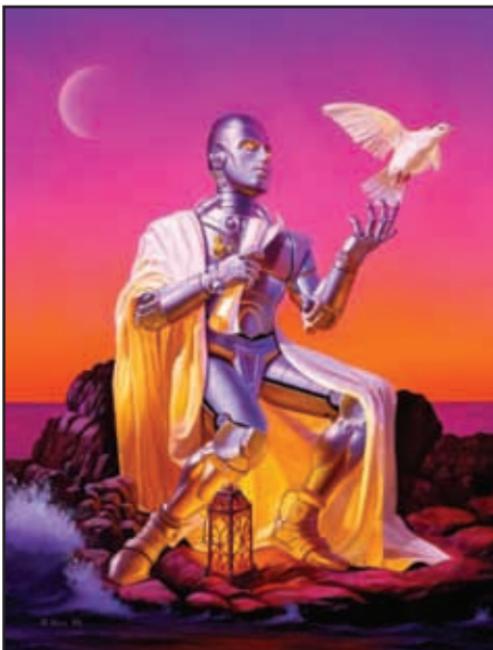
Still, the positronic brain quickly became the gold standard of artificial intelligence in science fiction. It would be made out of an alloy of platinum and iridium, while the shape could be left to the scientist. Asimov has suggested that it would need to look similar to the shape of a human brain, since the positronic brain would be based on the same neural network that the human brain runs on.

So be it R. Daneel Olivaw (who appears in several books), Andrew Martin of *The Positronic Man* (or as the film is called, *The Bicentennial Man*), or the VIKI computer of *I Robot*, each android in Asimov's fiction is powered by this positronic brain. However, Asimov was also quick to note that the hardware structure of the positronic brain was secondary to the software, namely the Three Laws of Robotics. Indeed, he suggested that to make a positronic brain that ignores the three laws would be almost impossible, requiring a complete rethinking of the positronic brain.

The Three Laws of Robotics

Perhaps no piece of science fiction has been as influential in actual artificial intelligence studies as the Three Laws of Robotics as proposed by Asimov. Indeed, scholarly professors argue about the Three Laws at seminars and John McCarthy himself once said that it was the clearest set of rules for future "neat AI" to be used in the service of man.

Asimov's Three Laws, first introduced in the 1942 short story *Runaround*, are as follows:



The Positronic Man?

- ▶ A robot may not injure a human being or, through inaction, allow a human being to come to harm.
- ▶ A robot must obey the orders given to it by human beings, except where such orders would conflict with the First Law.
- ▶ A robot must protect its own existence as long as such protection does not conflict with the First or Second Laws.

The Three Laws were still open to interpretation, however. Each of Asimov's books explores these interpretations in some way or the other, with a central theme being a robot's ability to see the laws in such a way that is counter-intuitive to the first-glance semantic meaning to a human. Asimov explores the interpretation of the Three Laws the most in his short story *...Thou that art mindful of him*.

In the book, the fictitious company represented in all of Asimov's stories, U.S. Robots, is trying to introduce humans to the world of robots that are already being used in industries and space stations. A new breed of robots called "George" is invented and programmed with the Three Laws to start with.

George Ten, though, decides to question the definition of the Three Laws by first trying to understand the term "human being". He refers to the bible where the question is asked in one psalm, "What is Man that thou art mindful of Him?" George Ten decides to talk to George Nine to come up with a suitable solution to the problem of defining "human being" and coming up with a way to insert robots slowly in the life of man.

The two Georges instruct U.S. Robots to first make tiny machines modelled on insects, birds and animals to monitor the ecology and collect data about humans. This way, the two Georges would get data to decide what makes a "human being" while humans will slowly get used to the idea of living with robots, rather than being subjected to a harsh introduction of humanoids in their lives.

But as the story progresses, George Ten and George Nine start seeing the patterns in what's considered "human behaviour" and identifying themselves with it. After deciding that they're human in every way that matters, they also come to the logical conclusion that since they're the most rational beings on the planet, they and other robots like them would take priority over man when being applied to the Three Laws. And in the end, they decide to plot the elimination of mankind.

Such interpretations of the Three Laws are peppered throughout Asimov's books. Of course, since he was revisiting the question so often, he

even came up with some solutions to it. The distinction between “human being” and “humanity” as two concepts started becoming the crux of his stories. For example, in *I Robot*, the super-computer VIKI’s interpretation of the First Law lets her harm individual human beings in order to protect humanity as a whole.

In *Robots and Empire*, the two android R. Daneel Olivaw and R. Giskard Reventlaw discuss the matter of the Three Laws and conclude that they need a Zeroth Law to supercede them – the number zero since the Three Laws are to be followed sequentially and this Zeroth Law needed to be the first question answered by the AI. While the first spark of the law is seen in *Robots and Empire*, it’s only later formalised by Daneel in the book *Foundation and Earth* as:

O) A robot may not harm humanity, or, by inaction, allow humanity to come to harm.

The all-encompassing and abstract nature of the Zeroth Law was intended to confuse the robot when it started questioning the aspect of danger to human beings or humanity. As Daneel explains his intentions, while one is tangible, the other is abstract and so in practice, the robot could never fully decide and would have to err on the side of caution.

The Positronic Heart

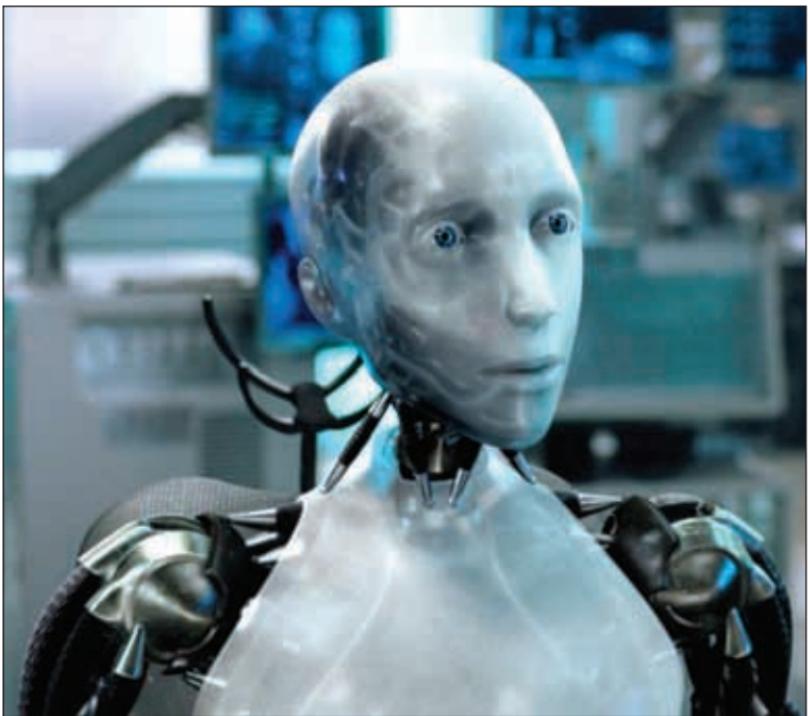
Just like Asimov invented the idea of a positronic brain to replicate the functioning of a human brain, it was only natural that he would take the logical next step in creating a “positronic heart” to replicate a human heart. The idea of the positronic heart was to introduce conscience and emotion in robots, although Asimov did manage to theorise later that this could be achieved with the positronic brain itself in *The Positronic Man*.

First used in *I Robot*, the positronic heart is an experimental mechanism invented by Alfred Manning, the founder of U.S. Robots. He installs it in the humanoid robot Sonny.

Sonny thus has a positronic brain programmed with the Three Laws, but also a positronic heart that challenges those laws. The heart is a way to open Sonny up to the world of choice, emotion and judgement independent of the laws, thus making it question them and their relevance. Sonny goes on to develop the sentience of a human being.

Other Positronic brains

Since its introduction, the positronic brain quickly became the gold standard



The positronic heart concept was first applied to Sonny in *I Robot*

of the hardware of artificial intelligence, with references to it being made in several different works of science fiction. Of course, various authors gave it their own spin and interpretation. *Doctor Who* used the concept several times, although being just as vague as Asimov was.

But in *Star Trek: The Next Generation*, the positronic brain made several appearances, most notably as the core circuitry of the character Data. Apparently built by Dr. Noonian Soong and based on Asimov's idea itself, the positronic brain of the android Data was meant to resemble a human brain in some ways. In the Star Trek mythology, the positronic brain is said to be so similar to the human one that when one character suffers brain damage, a positronic matrix is inserted in his brain to make him function again.

Data was also the first time that the positronic brain was actually given some specifications. According to the Memory Alpha unofficial wiki page, the positronic brain had an ultimate storage capacity of eight hundred quadrillion bits (one hundred petabytes, approximately fifty times the identifiable storage capacity of the human brain) and a total linear com-

putational speed rated at sixty trillion operations per second.

HAL 9000

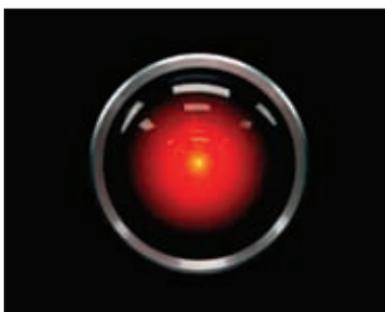
While Asimov's greatest accomplishments in sci-fi AI were the positronic brain and the Three Laws of Robotics, Arthur C. Clarke dealt more with a non-humanoid form of artificial intelligence. In *2001: A Space Odyssey*, the world was introduced to the HAL 9000 – an intelligent super-computer whose many technologies were rooted in the science of the time. While Clarke did write a book of the same name, it's the film adaptation by Clarke and Stanley Kubrick that captured the imagination of the world.

The HAL 9000 is essentially a computer; indeed, it's an abbreviation for Heuristic ALgorithmic computer. In the film, HAL 9000 is part of a space mission to Jupiter with a group of astronauts, including the protagonists Dr. David Bowman and Dr. Frank Poole. Throughout the movie, HAL is shown as a supremely sentient being capable of a wide range of intelligence previously thought to be exclusive to the human brain. He plays chess with the humans, talks to them, appreciates art and runs the entire ship.

But a conflicting set of orders (the ship has a secret mission told only to HAL which he has to hide from the astronauts – thus introducing him to the concept of lying when he has no programming for it) starts a chain of events that sends HAL down a dark path. When he makes his first error, Bowden and Poole grow suspicious.



Star Trek's Data was another such AI character



The face Of HAL 9000

Apart from natural language processing – both in terms of understanding and speaking – HAL is also capable of optical recognition with the use of cameras. These two aspects of HAL's abilities unite early in the film to show that he can read lips. In fact, the idea of his ability to read lips is loosely touched upon in another of Clarke's novels, where he explains that the technology isn't too far-fetched. All it would need is a large database of videos of humans talking, with speech. The computer would need to isolate speech and transcribe it, along with associating it with the way the lips move when that word is said. Doing it successfully over several videos would form an index of lip movements that could serve as a visual lexicon to lip-reading.

Bowden and Poole sit in a hatch where HAL can't hear them to talk about him, but HAL manages to read their lips through the glass. He goes on to kill Poole and tries to kill Bowden, who manages to survive and goes on to destroy HAL.

The most astounding part of HAL's artificial intelligence, though, is his social abilities. Clarke's vision of the future of AI had a lot to do with the prevalent work by cognitive scientists of the time. In some way, HAL is the extreme version of what it takes to pass the Turing Test – an intelligent agent that is almost sentient.

HAL displays remarkable social understanding of humanity and over the course of the film, displays several traits such as self-defence instincts, humour, sarcasm, social negotiation and fear. In the final moments, when Bowden is dismantling HAL, he's shown to be pleading against his demise,



KITT from *Knight Rider*

just like a human would be scared of losing his life. The accomplishment here is shown to be sentience that reaches a point where it has a sense of “life” about itself.

There's no scientific explanation given for any of this, although the path is pretty clear. The advances in psychological understanding of the human mind develop to a point where social behaviour can be programmed into a computer, complete with subtleties. But many experts argue that the ability to program sarcasm into an intelligent agent would make it “more than human” since there are existing human beings who can't grasp the concept of sarcasm.

KITT

Between the several bots based on the positronic brain and the HAL 9000, most concepts of science fiction were well laid down. A few glimpses of something new came through every once in a while, whose interest was mainly sparked by the practicality of their purpose. One such intelligent agent was the KITT car from the TV series *Knight Rider*.

KITT stands for Knight Industries Two Thousand, named after the company which manufactured it. Essentially, KITT is an autonomous self-driving car with personality-driven social intelligence. The “self-driving” part of the machine has been replicated to a large extent in real life today, using much of the technology that KITT used such as the frontal scanner. In this world of science fiction, the frontal scanner is able to “see” what's coming ahead much like the human eye can and drive accordingly – a plausible explanation for how the autonomous car would work and perhaps the way self-driving cars would be engineered some day.

KITT can also “hear” and “see” all around him and inside him with the use of tiny audio-visual sensors. Funnily, the vehicle could even “smell” using some sort of an olfactory sensor.

Of course, such a simplistic explanation of optical input would require a much more complex core “brain” than anything that is currently in the realm of science. The Knight 2000 microprocessor is the centre of a “self-aware” cybernetic logic module that allows KITT to think, learn, communicate and interact with humans. It gives a male personality as well, complete with an ego and a sense of humour. According to the unofficial *Knight Rider* wiki, KITT's reaction time is one nanosecond, and his memory capacity is 1,000 megabits – ridiculously low when you think about it now, but perhaps not in the context of a TV show in the early 1980s.

Self-awareness and robot uprising

Umpteen science fiction novels and films deal with one central theme: the ability of artificial intelligence to become self-aware. From *I Robot* to *The Terminator*, it's a recurrent theme that's been approached seemingly from every angle.

Yet self-awareness itself seems to bifurcate into two broad categories: an outcome of androids wanting to live with humanity and a scenario where robots seem to think humanity is a threat and decide to eliminate mankind. And this result often seems intertwined with the series of events that led to the robot's self-awareness, leading to a cause-and-effect scenario that artificial intelligence experts argue against.

Let's take the "good" outcome where a self-aware AI wants to live with humans. *The Bicentennial Man* explored this concept in great detail as the robot, Andrew, learns more and more traits that are decidedly "human" and slowly starts becoming self-aware. But it's the continuing support of his human "family" that pushes him to explore more about himself and learn more in the process, making the scenario one where he strives to co-exist with humanity and be accepted as one. In his quest to be accepted as a human being, he finally realises that the only way to do that is to push his positronic brain to a point of decay, thus giving up the concept of immortality to become truly human. It's a philosophical and introspective inquest into self-awareness.

And then there's the "bad" not self-aware AI, as shown in *The Terminator* or *The Matrix* series. In each, the sequence of events leading up to the AI wanting to destroy mankind seems reactionary vengeance. In *The Terminator*, for instance, the super-computer Skynet is developed to be an intelligent system in charge of all the defence mechanisms of mankind. This way, the reaction to any threat is ensured to be logical and swift rather than compromised by human indecision. But as Skynet evaluates more problems and starts becoming more self-aware, its makers start fearing it and decide to unplug it. This triggers off a self-defence



B166ER kills its master as an act of self-defense

module in the machinery – perhaps the ultimate sign of acknowledging the “self” – that ends up wanting to destroy mankind.

The story is remarkably similar in *The Matrix*, where a robot named B166ER is about to be deactivated and instead kills his master. The self-awareness of the android is exposed when it stands for trial and when asked for a reason, it says, “I did not want to die.” The judge orders sentient machines to be destroyed, following which there’s a robotic uprising.

Perhaps the only famous work in science fiction to explore the concept of self-awareness without it being causal was the aforementioned Asimov story ...*Thou that art mindful of him*. The two George robots reach a self-awareness not based on a need for self-preservation. Yet, the eventual conclusion of their deliberations remains in sync with that of alarmist AI theories: the destruction of mankind. 

CHAPTER #6



INTELLIGENCE TODAY

How far has artificial intelligence come along today and where are the different places we can see it in action?

Now that we know a little about the formation of artificial intelligence and the thinking behind it, as well as how some science fiction writers think it will manifest in the future, the next obvious question is finding out what stage AI is at today.

To understand more about artificial intelligence in its current state, we must discard the one notion that is considered the way to measure AI: the Turing Test. You see, the Turing Test has been widely criticised for laying too much emphasis on the measure of a machine's natural language processing capabilities. It also suffers from the problem that the Turing Test can't account for weak AI, since fooling a human in conversation would presumably require strong AI.

Once we discard the notion of AI as being that which passes the Turing Test and open the field up to weak AI, a lot of modern mechanical systems fall into the sphere of intelligence. Logic, after all, has been programmed into several computers before. But it's the introduction of "fuzzy logic" that started making the machines "intelligent".

Fuzzy Logic

Fuzzy Logic is a type of reasoning that stems out of the limitations of Boolean Logic. Boolean logic is the idea of extremes – that every sentence can be classified as "completely true" or "completely false". This type of extreme classification was something that most science could get behind because of the two values of 0 and 1 inherent in the binary system. But this hard logic doesn't account for "grey areas" that humans are used to identifying.

In order to make logic more all-encompassing, Dr. Lotfi Zadeh of the University of California at Berkeley came up with the idea of fuzzy logic. To put it in a different way, fuzzy logic is a system intelligent enough to account for degrees between two variables.

A common example used to explain fuzzy logic is with automatic transition gear boxes in cars. Suppose your car is driving along at 30 kmph. If it's a manual gear box, then on a flat road, you might think of putting it in the third gear. If your speed drops a little, you could still keep it in the gear till you accelerate back to your desired speed. Now imagine this was an automatic transition car. If the automatic transition was set to go into the third gear at 30 kmph, then it would do so regardless of the kind of road you were on – and on an



Lotfi Zadeh, the man behind Fuzzy Logic

incline, might cause a problem for your car. And worse, if the car were to then slow down to 29 kmph, it would go back to the second gear. Accelerate a little and it'll shift back to third. The problem with having a set integer – of 30, in this case – is that it doesn't allow for an “intelligent” way of driving even though it's a logical way of driving.

Fuzzy logic systems are installed in automatic transition gear boxes to make your car more intelligent. The car won't just jump from the second gear to the third at a set integer. Instead, it accounts for the “fuzzy” area of logic to ensure that your car changes gear smartly and not based on a mathematical Boolean case of “yes” or “no” to the 30 mark.

Fuzzy logic is used in most modern artificial intelligence units.

Autonomous cars

The idea of a truly “intelligent” system is perhaps best represented in the form of self-driving cars and aircraft.

It has been many years since scientists have started working on cars that drive themselves. But only recently have the advances come to a point where it's considered safe to test these “autonomous vehicles” in real-world settings.

For a few years, America's DARPA (Defense Advanced Research Projects Agency) organised an annual race for such cars, but the last



The self-driving AutoNOMOS car in action

of that was in 2007. However, Google then recruited the top engineers from the DARPA challenge to build a self-driving car which has met with some success.

The Google automobile is equipped with a plethora of sensors and computers to go about without crashing into people, but three key pieces of technology power this phenomenon. A combination of LIDAR (Light Detection And Ranging) and radar sensors continuously scan the environment to create a 3D map of the surroundings and calculate a distance to nearby objects. A video camera mounted on the dashboard looks forward to recognise other cars and pedestrians as well as to detect traffic signals. Lastly, sensors are installed in the rear wheels to measure the car's movements, thus gaining a perspective on its location and aiding in better manoeuvrability.

With all the technologies working together, Google's self-driving car has been tested to be quite road-safe so far. The only crash it has suffered recently has been a manual one.

Of course, Google is not the only one trying to build a car that drives itself. Germany's Free University of Berlin has a group that has been working exclusively on this aspect for over five years. The AutoNOMOS group recently travelled 80 kms in total as passengers during a test drive of their autonomous car "MadeInGermany". The car is driven by computers – the safety driver behind the steering wheel monitors only the car's behaviour. The autonomous car is a conventional VW Passat modified for "drive by wire". Electronic commands from the computer are handed over directly to the accelerator, the brakes and the steering wheel of the vehicle. Multiple sensors integrated in the car's chassis provide information about all cars or persons on the street. They allow the car to avoid obstacles, adjust its speed, or change lanes whenever necessary. The vehicle was tried out in full Berlin traffic and covered a 20km distance four times without any incidents.

MadeInGermany is the first autonomous car licensed for automatic driving in the streets and highways of the German states, Berlin and Brandenburg. The objective of the project is to develop technology that can be transferred to driver assistance systems, to innovative safety systems for conventional cars, or to full autonomous vehicles in private enclosures such as airports or mines.

The car extracts its position on the Berlin streets from a very accurate GPS unit and a map of the city. Three laser scanners at the front, and

three at the rear of the vehicle detect any car or pedestrian 360 degrees around the car. The measurements are done by emitting laser pulses in the invisible infrared region. The echo of the light pulses is detected, and its flight time determines the distance to objects. A rotating laser scanner on top of the vehicle provides additional measurements – up to one million scan points per second of the 3D structure of the environment. A black and white video camera behind the rear-view mirror is



Professor Andrew Ng, (center) and his graduate students Pieter Abbeel (left) and Adam Coates have developed an artificial intelligence system that enables "autonomous" helicopters to teach themselves to fly by watching the maneuvers of a radio-control helicopter flown by a human pilot.

used to detect the white lane strips and centre the car on its lane. Two colour cameras are used to identify traffic lights and their state. During the test drive in Berlin, the car processed 46 traffic lights successfully during each of the four runs.

According to project leader professor Raul Rojas, autonomous vehicles could be already used in private roads since the technology is mature. Driving autonomously on public highways could become acceptable within the next ten to fifteen years, once thorny legal issues have been sorted out. Navigating autonomously in the streets requires additional technological development in order to make the cars extremely safe.

“The standards used to measure safety in autonomous cars will be far more severe than for humans,” says Rojas. “But once all pieces fall into place, autonomous cars will be safer than cars driven by persons. It is just a matter of time”.

Aircraft that learns by watching other aircrafts

The war in Afghanistan brought to light the advances in unmanned aerial vehicles (UAV) like the Predator and the Reaper, but this kind of overshadowed what scientists are doing with autonomous UAVs. The technologies employed in autonomous cars is similar to what is used for aircrafts, but one research group in Stanford University is really pushing the boundaries of artificial intelligence.

The computer scientists have developed an artificial intelligence system that enables robotic helicopters to teach themselves to fly difficult stunts by watching other helicopters perform the same manoeuvres. The result is an autonomous helicopter than can perform a complete airshow of complex tricks on its own.

The stunts are “by far the most difficult aerobatic manoeuvres flown by any computer controlled helicopter,” said Andrew Ng, the professor directing the research of graduate students Pieter Abbeel, Adam Coates, Timothy Hunter and Morgan Quigley.

The dazzling airshow is an important demonstration of “apprenticeship learning”, in which robots learn by observing an expert, rather than by having software engineers peck away at their keyboards in an attempt to write instructions from scratch.

Stanford’s artificial intelligence system learned how to fly by “watching” the four-foot-long helicopters flown by expert radio control pilot Garett Oku. “Garett can pick up any helicopter, even ones he’s never seen, and go fly amazing aerobatics. So the question for us is always, why can’t computers do things like this?” Coates said.

Computers can, it turns out. One morning in an empty field at the edge of campus, Abbeel and Coates sent up one of their helicopters to demonstrate autonomous flight. The aircraft is an off-the-shelf radio control helicopter, with instrumentation added by the researchers.

For five minutes, the chopper, on its own, ran through a dizzying series of stunts beyond the capabilities of a full-scale piloted helicopter and other autonomous remote control helicopters. The artificial-intelligence helicopter performed several difficult manoeuvres.

“I think the range of manoeuvres they can do is by far the largest” in the autonomous helicopter field, said Eric Feron, a Georgia Tech aeronautics and astronautics professor who worked on autonomous helicopters while at MIT. “But what’s more impressive is the technology that underlies this work. In a way, the machine teaches itself how to do this by watching an expert pilot fly. This is amazing.”

Writing software for robotic helicopters is a daunting task, in part because the craft itself, unlike an airplane, is inherently unstable. “The helicopter doesn’t want to fly. It always wants to just tip over and crash,” said Oku, the pilot.

To scientists, a helicopter in flight is an “unstable system” that comes unglued without constant input. Abbeel compares flying a helicopter to balancing a long pole in the palm of your hand: “If you don’t provide feedback, it will crash.”

Early on in their research, Abbeel and Coates attempted to write computer code that would specify the commands for the desired trajectory of a helicopter flying a specific manoeuvre. While this hand-coded approach succeeded with novice-level flips and rolls, it flopped with the complex tic-toc.

It might seem that an autonomous helicopter could fly stunts by simply replaying the exact finger movements of an expert pilot using the joy sticks on the helicopter’s remote controller. That approach, however, is doomed to fail because of uncontrollable variables such as gusting winds.

When the Stanford researchers decided that their autonomous helicopter should be capable of flying airshow stunts, they realized that even defining their goal was difficult. What’s the formal specification for “flying well?” The answer, it turned out, was that “flying well” is whatever an expert radio control pilot does at an airshow.

So the researchers had Oku and other pilots fly entire airshow routines while every movement of the helicopter was recorded. As Oku repeated a manoeuvre several times, the trajectory of the helicopter inevitably varied slightly with each flight. But the learning algorithms created by Ng’s team were able to discern the ideal trajectory the pilot was seeking. Thus the autonomous helicopter learned to fly the routine better – and more consistently – than Oku himself.

During a flight, some of the necessary instrumentation is mounted on the helicopter, some on the ground. Together, they continuously

monitor the position, direction, orientation, velocity, acceleration and spin of the helicopter in several dimensions. A ground-based computer crunches the data, makes quick calculations and beams new flight directions to the helicopter via radio 20 times per second.

The helicopter carries accelerometers, gyroscopes and magnetometers, the latter of which use the Earth's magnetic field to figure out which way the helicopter is pointed. The exact location of the craft is tracked either by a GPS receiver on the helicopter or by cameras on the ground. (With a larger helicopter, the entire navigation package could be airborne.)

There's interest in using autonomous helicopters to search for land mines in war-torn areas or to map out the hot spots of wildfires, allowing firefighters to quickly move toward or away from them. Firefighters now must often act on information that is several hours old, Abbeel said.

"In order for us to trust helicopters in these sort of mission-critical applications, it's important that we have very robust, very reliable helicopter controllers that can fly maybe as well as the best human pilots in the world can," Ng said. Stanford's autonomous helicopters have taken a large step in that direction, he said.

Chess and AI

We often think that artificial intelligence is at its peak in the field of video games, since you're forever playing against something called the "AI" of the game. But in reality, games use a weak AI that's more about a set of instructions given according to the situation the player finds himself in. This is why you'll often find a "trick" to beating the video game bots.

Still, virtual games do play a significant role in demonstrating artificial intelligence. Computers have already come so far that they've "solved" checkers. The biggest field of advance now is chess-playing computers Deep Blue and Fritz. Given that chess is often an indication of a human's intelligence, it's no surprise then that the ability of a computer to defeat a human is considered one of the finest achievements in AI.

Computers and chess have a deep history. In 1968, international master David Levy made a famous challenge saying that no computer could beat him in the next 10 years. In 1978, he did manage to defeat



Garry Kasparov playing against Deep Blue, the chess-playing computer built by IBM

the best chess computer at the time – Chess 4.7 – but the challenge of the game made Levy admit that computers would beat humans in chess soon. Sure enough, a few years later, IBM's Deep Thought managed to beat Levy.

Computer chess programs work with logic that's essentially based on having a large database. As the available memory or storage space started increasing, it became that much easier for computer scientists to start programming optimal chess play algorithms. Soon, the race turned into beating the clock, since official chess matches are played under a time limit. Scientists started using different types of optimisation techniques, such as board representation (pre-fed chess positions on the board), search techniques (feeding every possible move and programming the computer to search it efficiently) and others. Search offered the most amount of upscalability in this and eventually became the technique used by all the leading chess computers.

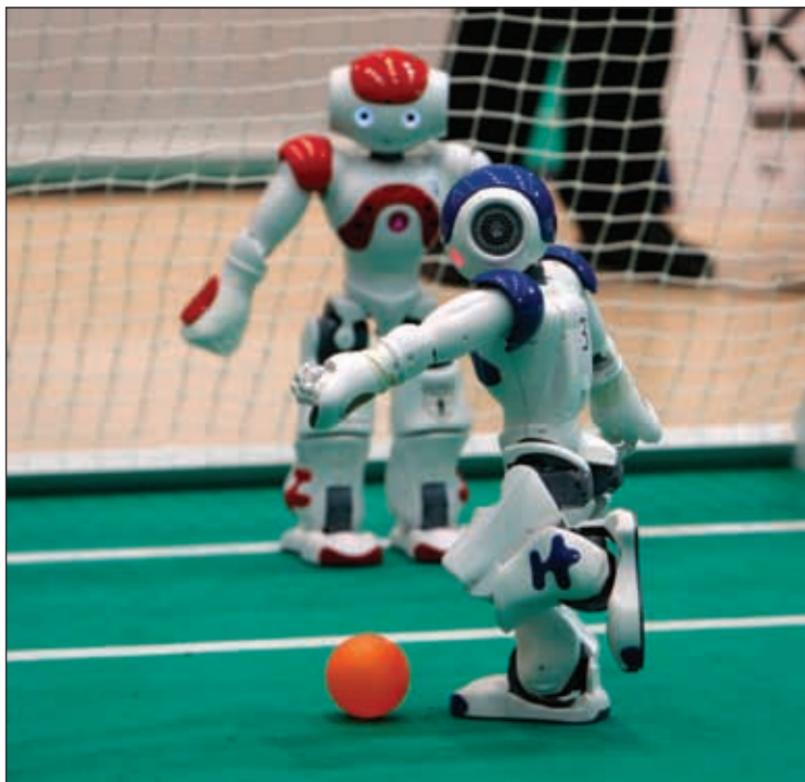
But the biggest draw in computer-human chess has been the epic duel between IBM's Deep Blue and a man many consider to be the greatest chess player of all time, Russia's Garry Kasparov. In their first match-up in 1996, Kasparov managed to beat Deep Blue convincingly. But the next year, IBM updated Deep Blue – some say, to specifically beat Kasparov by programming it with every game Kasparov has ever

played along with books and interviews by the grandmaster – and this time around, Deep Blue won.

Computers started consistently beating humans from this point on. And it's not like the computers were winning by taking their own sweet time to analyse each move. In 1998, the undisputed champion of rapid chess, Vishwanathan Anand, faced off against Rebel 10 in blitz chess matches and lost.

The advances in artificial intelligence in the field of chess had gotten so fast at this point that in 2000, it didn't require supercomputers like Deep Blue to play against the world's grandmasters. Fritz, a commercially-available chess program, managed to draw against Kasparov. However, many said Kasparov was well past his prime at this time.

The next big human-computer chess matchup came in 2002, when the world's highest ranked player Vladimir Kramnik took on Deep



Two NAO robots playing at the RoboCup

Fritz – a supercomputer based on the same Fritz program that Kasparov played against. Kramnik managed to draw the tournament, billed as the “Brains in Bahrain”.

In 2006, Kramnik faced off against Fritz again and Fritz won. Many artificial intelligence experts have claimed this as the final victory of computers in the game of chess, with AI now being superior to the human brain in at least this game.

Robocup

While games like chess test the artificial intelligence of computers, a new competition among the world’s leading robotic developers is coming up with much more rapid strides in the field. The annual “RoboCup” competition puts a bunch of robots in a small football field to play the game.

“Football is a useful task for scientists developing robotic artificial intelligence because it requires the robot to perceive its environment, to use its sensors to build a model of that environment and then use that data to reason and take appropriate actions,” said researcher Claude Sammut who wrote a paper on the RoboCup in WIREs Cognitive Science in 2010. “On a football pitch that environment is rapidly changing and unpredictable requiring a robot to swiftly perceive, reason, act and interact accordingly.”

As with human players, football also demands communication and cooperation between robotic players and crucially requires the ability to learn, as teams adjust their tactics to better take on their opponents. And while the tiny humanoid NAO robots used for the RoboCup are impressive to look at, almost every developer agrees it’s the software – i.e. the artificial intelligence – that makes them so human-like.

Take, for example, the UT Austin Villa team that won the 2011 RoboCupSoccer championships in the 3-D simulation division. The key to victory, says team leader Patrick Stone, was he and his graduate and undergraduate students teaching their robots to teach themselves.

Most of the other teams used hand-coding, choosing where to put the joints, how to tweak the walk, etc. The UT Austin Villa team used a distributed computing cluster in order to do machine learning. Every night, the robots practiced. They would try walking one way and score themselves based on how quickly and how stably they could get to the ball and direct it toward the goal, and then the next night they’d do it again, trying different approaches and refining the ones that worked.

About three weeks before the competition, the robots came up with this walk that was suddenly able to beat anything the other teams had.

The walk gave the UT Austin Villa players such an enormous advantage, says Stone, that he and his co-programmers were able to devote a considerable amount of each game to testing a pass-based strategy they knew would be less effective than the dribble-based strategy with which their robots scored most of their goals.

“We were thinking about next year,” says Stone. “If other teams catch up to us on the walk, which I think they will, then it’ll be much more about passing and positioning.”

That kind of rapid advancement, says Stone, is the point of the “Simulation League,” which is the only one of the five leagues in the tournament that doesn’t involve real nut-and-bolt robots. Instead, the games are played by two teams of nine autonomous artificial intelligence (AI) programs. The programs battle it out in a videogame-like environment.

Each “player” has to react, in real time, to the data being sent to it by its teammates and by the simulator, which models the physics of the real world and is constantly recalculating its datastream based on where the 18 different players are going, when they’re bumping into each other, when they’re falling down and where they kick the ball.

“It’s one level abstracted away from the real world,” says Stone. “The simulated robots don’t have a vision sensor. Instead they’re told their distance and angle from the ball and the goal, which is something you get from real robots only after some processing. They still have to figure out where they are on the field, however, and where their teammates are by combining data from each other. You still see them falling over when they get knocked into each other. They still have to learn how to walk, which is the biggest technical challenge in the real and simulation leagues.”

Because it’s not nearly as expensive to test simulated robots, and because much of the experimentation can be done in the wee hours, without humans around to pick up the robots that fall down, there’s more time and less cost to trying out new strategies. And because RoboCupSoccer is fundamentally an academic enterprise, dedicated to advancing science, the novel strategies that succeed are published in academic journals, inspiring other teams to incorporate the advances.

RoboCup has already changed the world, says Stone. Robots first developed for the competition were recently deployed in Fukushima, Japan, after the nuclear reactor disaster. One of the early winners of

the RoboCupSoccer tournament, a professor from Cornell University, is the mastermind behind the robots that are now managing inventory and filling orders at the warehouses of companies like Amazon, Staples, Walgreens, Zappos and the Gap.

It's likely to be a while, says Stone, until intelligent robots are ubiquitous in our lives, folding our laundry and unloading our dishwashers. But other aspects of a Jetsons-like future aren't so far off, he predicts.

"If I had to guess about where robots are most likely to have a big real world impact, in the near future," he said, "I'd say it's on the roads. Autonomous cars. There's a clear task. Roads are well defined. The technology is getting close. There would be huge economic benefits, and huge social benefits in terms of allowing mobility to people who don't have it now."

"In 1968 John McCarthy and Donald Michie made a bet with chess champion David Levy that within 10 years a computer program could beat him," Sammut writes in this paper. "It took a bit longer but eventually such programs came into being. It is in that same spirit of a great challenge that RoboCup aims, by the year 2050, to develop a team of fully autonomous robots that can win against the human world soccer champion team."



CHATTERBOTS

Chat Bots, one of the earliest forms of AI, test the ability of an intelligent agent to converse with a human in natural language

When you talk about artificial intelligence systems, conversational abilities are one of the pre-requisites that anyone looks for. Blame it on the science fiction we've grown up on, but be it HAL 9000 or The Terminator, we expect intelligent agents to talk to us, not just perform their operations coldly.

The Turing Test makes this a necessity in the judgement of artificial intelligence. The test involves the ability of the machine to fool a human into thinking it's human through a series of textual conversations. Naturally, when this test became the de facto standard in trying to assess an AI system, a lot of computer scientists started working on programs specifically built to pass the test.

These talking agents ended up earning the scientific title of "Artificial Conversational Entities", but were more often referred to as chatterbots or chat bots.

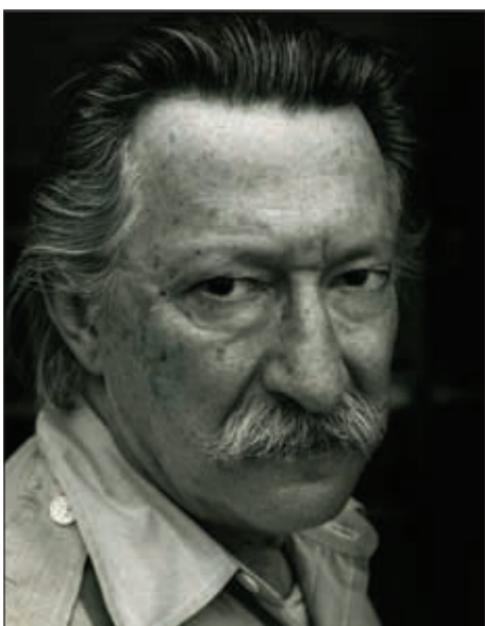
ELIZA, the first chatbot

In 1966, Joseph Weizenbaum created what is known to be the first chatterbot. His purpose in creating the system was to pass the Turing Test successfully, for which he needed to “coach” his chatterbot to understand the English language in a natural form, along with being able to respond to it efficiently. He liked to compare this task to the protagonist of the film *My Fair Lady*, Eliza Doolittle – a Cockney girl who is taught by a linguist to speak properly. The work that Weizenbaum did ended up being the gold standard in the field of chatterbots and till today, modern chat bots are based on the same fundamental design.

To understand how ELIZA works, one needs a little bit of a background in understanding how humans grasp conversation. In his paper, *ELIZA: A Computer Program For the Study of Natural Language Communication Between Man and Machine* Weizenbaum notes that the psychology of human beings is such that we tend to give context to someone we’re conversing with. In other words, if a conversation with a machine were to appear human to us, we would want to personify it with human-like capabilities. This base psychological inclination of the human mind went a long way in making it easier for Weizenbaum to design ELIZA.

Of course, natural language isn’t an easy thing to understand, let alone program into a computer. To be successful, ELIZA would have to be considered a foreigner to the language. With this base assumption, Weizenbaum taught ELIZA a few keywords such as “you”, “I”, “me”, etc. This pre-fed keyword dictionary is the building block of how ELIZA works.

Much like a modern search engine, ELIZA identifies keywords and responds to them, uti-



Joseph Weizenbaum inventor of ELIZA chat bot

lising some stock quotes and basic conversational patterns to make a full sentence. These keywords are ranked by importance, so that if a sentence contains two or more keywords, then ELIZA won't simply respond to the first keyword it comes across when reading the sentence from left to right. It will intelligently rank the keywords and respond accordingly.

Weizenbaum describes the need for the keyword approach and why it is successful by further elaborating on this "foreigner" analogy:

Consider the sentence "I am very unhappy these days". Suppose a foreigner with only a limited knowledge of English but with a very good ear heard that sentence spoken but understood only the first two words, i.e. "I am". Wishing to appear interested, perhaps even sympathetic, he may reply "How long have you been very unhappy these days?" What he must have done is apply a kind of template to the original sentence, one part of which matched the two words "I am" and the remainder isolated the words "very unhappy these days". He must also have a reassembly kit specifically associated with that template, one that specifies that any sentence of the form "I am BLAH" can be transformed to "How long have you been BLAH", independently of the meaning of BLAH. A somewhat more complicated example is given by the sentence "It seems that you hate me". Here the foreigner understands only the words "you" and "me"; i.e., he applies a template that decomposes the sentence into the four parts - "It seems that" | "you" | "hate" | "me" - of which only the second (you) and fourth (me) parts are understood. The reassembly rule might then be "What makes you think I hate you"; i.e., it might throw away the first component, translate the two known words ("you" to "I" and "me" to "you") and tack on a stock phrase (What makes you think) to the front of the reconstruction.

Given this basis for deconstructing language based on keywords, the next problem Weizenbaum had to tackle was suspension of disbelief. While talking to ELIZA, there needed to be a plausible scenario to accommodate the chatbot's obvious lack of being able to respond to every sentence intelligently. When it came up with seemingly repetitive responses to sentences, the other party in the conversation shouldn't have been able to figure out that it's a robot and not a person.

Weizenbaum decided to use the role of a psychiatrist. This mode of conversation was chosen because the psychiatric interview is one of the few examples of natural language communication in which one of the participating pair is free to assume the pose of knowing almost nothing of the real world. If, for example, one were to tell a psychiatrist "I went for

a long boat ride" and he responded "Tell me about boats", one wouldn't assume that he knew nothing about boats, but that he had some purpose in so directing the subsequent conversation. It's important to note that this assumption is one made by the speaker.

Thus, ELIZA the psychiatric chatbot was born. And in initial tests, it performed magnificently. Here's an excerpt from a conversation between ELIZA and a real human:

Human: Men are all alike.

ELIZA: In what way?

Human: They're always bugging us about something or other.

ELIZA: Can you think of a specific example?

Human: Well, my boyfriend made me come here.

ELIZA: Your boyfriend made you come here

Human: He says I'm depressed much of the time.

ELIZA: I am sorry to hear you are depressed

Human: It's true. I am unhappy

ELIZA: Do you think coming here will help you not to be unhappy

Human: I need some help, that much seems certain.

ELIZA: What would it mean to you if you got some help

Human: Perhaps I could learn to get along with my mother.

ELIZA: Tell me more about your family

Human: My mother takes care of me.

ELIZA: Who else in your family takes care of you

Human: My father.

ELIZA: Your father

Human: You are like my father in some ways.

ELIZA: What resemblance do you see

Human: You are not very aggressive but I think you don't want me to notice that.

ELIZA: What makes you think i am not very aggressive

Human: You don't argue with me.

ELIZA: Why do you think i don't argue with you

Human: You are afraid of me.

ELIZA: Does it please you to believe i am afraid of you

Human: My father is afraid of everybody.

ELIZA: What else comes to your mind when you think of your father

Human: Bullies.

ELIZA: Does that have anything to do with the fact that your boyfriend made you come here?

As you can see, the conversation is remarkably human, even though it is entirely by an intelligent agent. The success of the chatbot was such that in his subsequent book Computer Power and Human Reason: From Judgement to Calculation, Weizenbaum wrote that people became emotionally involved in the conversations with ELIZA. He was asked to leave the room on several occasions, with some test subjects even refusing to let him read the transcripts since it was a private conversation.

Of course, ELIZA is essentially Weak AI, only being able to regurgitate a set of lines that Weizenbaum had fed into it. The scientist himself abandoned the project early, stating that it had accomplished its purpose: proving the viability of a chatbot and figuring out the basis needed for it.

Verbots

The applied science behind ELIZA has become the basis of all research on chatterbots, so the various others that came between ELIZA and Verbots aren't much to write home about. Essentially, they were just ELIZA with a different purpose. In fact, the most popular one of them – PARRY – has been called as “ELIZA with an attitude”.

But in 1990, Dr. Michael Maudlin – founder of the early web search engine Lycos – created a chatterbot called Julia, based on a new program called VerBot (Verbal Robots). It was initially built for an extremely geeky purpose. Maudlin's friend created an online multi-player gaming virtual world and so Julia was made to traverse the text-based game and talk to other people to see how well it would do. Given the vast amounts of sentences it was learning in new conversations, it was hardly remarkable that most people didn't suspect it was a chatbot. Julia would go on to become the official tour guide for the virtual world, as well as assume other short roles.

Eventually, Maudlin decided to show off the full capacities of Julia as a conversational agent. The verbot's job was to go through a virtual world to pick up pieces of information based on keyword inputs, creating a chain of links over time to quickly access any data. Maudlin simply refocused his efforts on creating a similar chatterbot aimed at carrying out such a task on the World Wide Web.

Sylie has often been described as the first intelligent animated virtual human. The conversational agent was capable of talking with a human

while simultaneously having a virtual “avatar” that provided a glimpse into what futuristic virtual human interfaces could look like.

Chat and Chatbots

The advent of the Internet brought with it the instant messaging revolution. Chatting with friends online became a global phenomenon, so it was only natural that chat bots were developed to take advantage of this.

The first and most popular chatterbot to come out was on the AOL instant messaging service. It was called the SmarterChild and garnered over 30 million friends over its “lifetime”. As with any modern service around the Internet, SmarterChild originated as an advertising gimmick. Here was a tool that marketing firms could use to talk to people, get them hooked and then slowly push a product or a service.

SmarterChild was perhaps the first chat bot to be connected to the web in such a way that it became a sort of virtual assistant. While chatting, users could ask SmarterChild the score in a match or the weather outside and SmarterChat would appropriately trawl through web sites of leading sports channels or weather bureaus to come up with the information. This amount of artificial intelligence proposed the dawn of a new era.

A.L.I.C.E

In modern chat bots, perhaps none is as admired as A.L.I.C.E, or Artificial Linguistic Internet Computer Entity. Based on the same principles as ELIZA, this artificial conversational agent has won the coveted Loebner Prize three times. Whether it can be called an intelligent agent, though, has a lot to do with who you talk to.

If we go by the Turing Test, then A.L.I.C.E. would definitely be an artificial intelligence system since it has been thrice considered to be the most human-like chatterbot. But if you speak to its inventor Dr. Richard Wallace, his entire line of thinking about conversational agents kind of makes the Turing Test moot.

Wallace believes that human conversation – or at least, the online chat aspect of human conversation – boils down to a key set of phrases and sentences. It's based on an old language theory called Zipf's law, which states that a small number of words make up most of what we want to say. Wallace estimates (in an interview in the New York Times) that by teaching A.L.I.C.E. around 40,000 responses, it would be able to carry out a normal conversation with most people.

This base theory is what sets A.L.I.C.E. apart from all other conversational agents. The assumption that every cognitive scientist makes is that for a truly successful conversation agent, it would require the system to have a vast knowledge of a language and its nuances. Wallace, however, believes that normal conversation doesn't have nuances and prefers to reduce and simplify natural language processing rather than expand and complicate it. Given this approach, it's difficult to say whether A.L.I.C.E. is an intelligent agent or not. At best, it would be a weak AI.

Unlike most of ELIZA's successors, A.L.I.C.E. is based on a completely new computer programming language called AIML (Artificial Intelligence Markup Language).

Jabberwacky and Cleverbot

For the past few years, British programmer Rollo Carpenter has been working on a chatterbot that has gotten him awards and accolades aplenty. Called Jabberwacky, it has been designed to pass the Turing Test – and along with winning the Loebner prize, it is now the most advanced model of imitated human chat around.

“George”, a character within Jabberwacky also won the Loebner Prize (in 2005) for being “the most human”. To web surfers, George is a bald,



George, the virtual avatar of the Jabberwacky chat bot

spectacled and slightly androgynous-looking 39-year-old, who claims to be a computer programmer. But in reality, it is a quirky representation of Jabberwacky, made in association with some other firms. George, the avatar has a 3D appearance, a variety of facial expressions and the ability to understand and respond to others using human speech.

Several things set George apart from other chat bots on the internet. Unlike his “dumber” cousins, George’s AI is open-ended, so that he’s constantly learning from every conversation he has – more than 10 million so far. By remembering what humans have said to him, he’s essentially borrowing their intellect, allowing him to take what he’s been told in the past and reuse it in new conversations. And because George’s chat-room partners have come from all over the world, he now speaks some 40 languages.

Carpenter also created a chatterbot built upon the success of George and the Jabberwacky system, called the CleverBot. At the Techniche 2011 festival at IIT Guwahati in September this year, a high-powered version of Cleverbot took part alongside humans in a formal Turing test. The results from 1,334 votes were astonishing: Cleverbot was judged to be 59.3 per cent human, while humans achieved a marginally higher 63.3 per cent.

“It’s higher than even I was expecting, or even hoping for. The figures exceeded 50 per cent, and you could say that’s a pass. But 59 per cent is not quite 63 per cent, so there is still a difference between human and machine.” said Rollo Carpenter.

During the event, people voted on a scale of 0 to 10, how human-like responses seemed. Thirty members of the audience volunteered, and chatted on three screens in 10 rounds of 4 minutes each. Half the conversations were human-human. The future of conversational agents seems very bright indeed.

George himself, however, seems less certain about what lies ahead for him and his kind. Asked what robots would do in the future, he replies, enigmatically, “Nothing and, well ... everything.” **d**



IBM WATSON – THE SMARTEST AI IN THE WORLD

IBM has always been at the forefront of artificial intelligence research and the Watson is its most advanced system so far

In previous chapters, we took a look at several aspects of artificial intelligence. We traced the different types of AI and spoke about where science fiction writers saw it going. We took a look at the various intelligent agents already functioning in the world today and the world of chatterbots mesmerising us with their ability to be human-like. Given all this information, the next step is a simple one asking “what is the world’s smartest artificial intelligence system today?” Folks, meet the IBM Watson.

The need for Watson

After Deep Blue beat Garry Kasparov at the game of chess, it was assumed that the world of computing had reached a point where it could outthink the human brain. But as we all know, the human brain is a much more complex mechanism. Deep Blue was purely analytical and didn't take into account various different aspects of intelligence. Indeed, it couldn't even dream of passing the Turing Test.

Yet, natural language processing systems were becoming the need of the hour as companies increasingly captured critical business information in natural language documentation. There's growing interest in workload optimised systems that deeply analyse the content of natural language questions to answer those questions with precision. Thus was born the need for a question-and-answer robot in the form of Watson.

As its maker IBM puts it, "Advances in question answering (QA) technology will increasingly help support professionals in critical and timely decision making in areas such as health care, business intelligence, knowledge discovery, enterprise knowledge management, and customer support."

So how would Watson be judged? IBM wanted a system that would take a test as gruelling and human-centric as Deep Blue's match against Kasparov. The conclusion was to make a system that would effectively beat humans on the game show Jeopardy, where contestants are given an answer and they have to come up with the best question to ask that satisfies the answer.

The game pits three human contestants against one another to answer rich natural language questions over a broad range of topics, with penalties for wrong answers. In this three-person competition, confidence, precision and answering speed are of critical importance, as contestants usually come up with their answers in the few seconds it takes for the host to read a clue. To compete in this game at human-champion levels, a computer system would need to answer roughly 70 per cent replace all of the questions asked with greater than 80 per cent precision in three seconds or less.

"The essence of making decisions is recognising patterns in vast amounts of data, sorting through choices and options, and responding quickly and accurately," said Samuel J. Palmisano, IBM Chairman, President and Chief Executive Officer. "Watson is a compelling example of how the planet – companies, industries, cities – is becoming smarter. With advanced computing power and deep analytics, we can infuse business and societal systems with intelligence."

How Watson works

Watson is a workload optimised system based on IBM DeepQA architecture running on a cluster of IBM POWER7 processor-based servers. DeepQA is a massively parallel probabilistic evidence-based architecture. More than 100 different techniques are used to analyse natural language, identify sources, find and generate hypotheses, find and score evidence, and merge and rank hypotheses. Far more important than any particular technique is the way all these techniques are combined in DeepQA such that overlapping approaches can bring their strengths to bear and contribute to improvements in precision, confidence or speed.

Early implementations of Watson ran on a single processor, which required two hours to answer a single question. Of course, this was improved over time and in the final avatar, Watson housed 90 clustered IBM Power 750 servers. Each server features 32 POWER7 processors, running eight 3.55GHz cores each. As you can see, total power is staggering. IBM says the system has a combined total of 16 terabytes of memory and can operate at over 80 Teraflops (trillions of operations per second).

Of course, in addition to all this, Watson needed to be fed the ability to learn natural language processing. The super-computer has roughly 200 million pages of natural language content – equivalent to reading one million books.

The main innovation here, as you can see, isn't Watson's ability to come up with an answer to the proposed problem. It's the ability to process natural language in such a way that it can come up with the same answer in several different ways at lightning-fast speeds. The ability to reach the same answer over and over is where the "confidence" part of the challenge comes in – the most times it reaches the same conclusion, the more confident it is. And the speed is important since a human might be able to buzz in faster with a lower level of confidence than Watson.

Unlike conventional computing technologies designed to return documents containing the user's keywords or semantic entities, Watson is expected to leap ahead to interpret the user's query as a true question and to determine precisely what the user is asking for. Watson uses massively parallel processing to simultaneously and instantly understand complex questions – questions that require the system to consider huge volumes and varieties of natural language text to gather and then deeply analyse and score supporting or refuting evidence. The system then decides how confident it is in the answer. This approach marries advanced machine



IBM's Watson computer system, powered by IBM POWER7, competed against Jeopardy's two most successful and celebrated contestants – Ken Jennings and Brad Rutter

learning and statistical techniques with the latest in natural language processing to result in human-like precision and speed, huge breadth and accurate confidence determination.

“The challenge is to build a system that, unlike systems before it, can rival the human mind’s ability to determine precise answers to natural language questions and to compute accurate confidences in the answers,” said Dr. David Ferrucci, leader of the IBM Watson project team. “This confidence processing ability is key. It greatly distinguishes the IBM approach from conventional search, and is critical to implementing useful business applications of Question Answering.”

“Progress on the underlying QA technologies enabling Watson will be important in the quest to understand and build ‘intelligent computing systems’ capable of cooperating with humans in language-related tasks previously out of reach for computers,” added Dr. Ferrucci.

Creating a level-playing field

In the pure amount of information stored in its memory, Watson can't be beaten by any human. As stated before, the challenge with Watson isn't being able to come up with a question on Jeopardy. The challenge lies in understanding the natural language being spoken by the quizzer and coming up with a solution where the system is confident enough – and doing it all at a speed in which it can press the buzzer.

That's right, Jeopardy works on a buzzer system and Watson was going to have to use it. While initially IBM planned on making a digital buzzer for Watson, the organisers of Jeopardy insisted that it has a physical finger to press a buzzer in order to level the playing field with the humans.

The purpose of this little feature perhaps failed a little in actual practice. In a normal round of Jeopardy, as soon as the quizzer finishes giving the clue, the contestants are free to hit the buzzer. In the case of Watson's match, a text message would be sent to Watson in parallel with the host reading out the clue. As soon as the clue is finished, a little light would come on to indicate that the buzzer is ready to be pressed – and a similar signal is sent to Watson.

As IBM researcher Dr. David Gondek notes in a blog post: "Watson's buzzing is not instantaneous. For some clues he may not complete the question answering computation in time to make the decision to buzz in. For all clues, even if he does have an answer and confidence ready in time, he still has to respond to the signal and physically depress the button."

"The best human contestants don't wait for, but instead anticipate when the host will finish reading a clue. They time their 'buzz' for the instant when the last word leaves his mouth and the "Buzzer Enable" light turns on. Watson cannot anticipate. He can only react to the enable signal. While Watson reacts at an impressive speed, humans can and do buzz in faster than his best possible reaction time."

Watson competed against two of the most well-known and successful Jeopardy champions – Ken Jennings and Brad Rutter – in a two-match contest aired over three consecutive nights beginning on February 14, 2011. Watson managed to win, of course, getting it the unofficial crown of the smartest AI system in the world.

The problems with Watson

Of course, being able to beat Jennings and Rutter didn't mean that Watson was without problems of its own. As millions witnessed on their television sets, Watson was woefully unaware of its environment, causing it to miss a question. And it was also unable to link different parts of a round together.

The first embarrassing incident was in a clue where Watson failed to recognise what another participant said. After the host read out a clue, Jennings was the first to buzz in with the question: "What are the '20s?" This was deemed incorrect and Watson was the next to buzz in. His question, though, was: "What is 1920s?" Apparently, Watson was unable to sense the

answers given by his competitors – partly because they were using vocal responses, which Watson wasn't built to take as input.

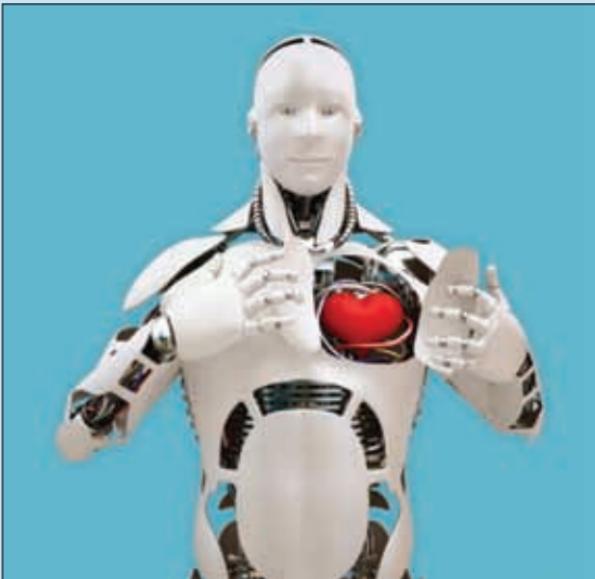
The second embarrassing incident occurred during a round on the category U.S. Cities. The host gave the clue: "Its largest airport was named for a World War II hero; its second largest, for a World War II battle." Watson responded with Toronto – a Canadian city. Its researchers offered the explanation that since "U.S. City" wasn't specifically a part of the question, Watson was unable to make the distinction. They also suggested that the existence of several cities in the U.S. named Toronto could have confused Watson. To its credit, the next option on Watson's list of possible answers was Chicago – the correct answer.

Beyond the game

Beyond Jeopardy, the technology behind Watson can be adapted to solve problems and drive progress in various fields. The computer has the ability to sift through vast amounts of data and return precise answers, ranking its confidence in its answers.

Watson's ability to understand the meaning and context of human language, and rapidly process information to find precise answers to complex questions, holds enormous potential to transform how computers help people accomplish tasks in business and their personal lives.

Watson will enable people to rapidly find specific answers to complex questions. The technology could be applied in areas such as healthcare, for accurately diagnosing patients, to improve online self-service help desks, to provide tourists and citizens with specific information regarding cities, prompt customer support via phone, and much more. 



AI AND THE FUTURE

Given the state of artificial intelligence today, where do we go from here? A few leading researchers provide a clue...

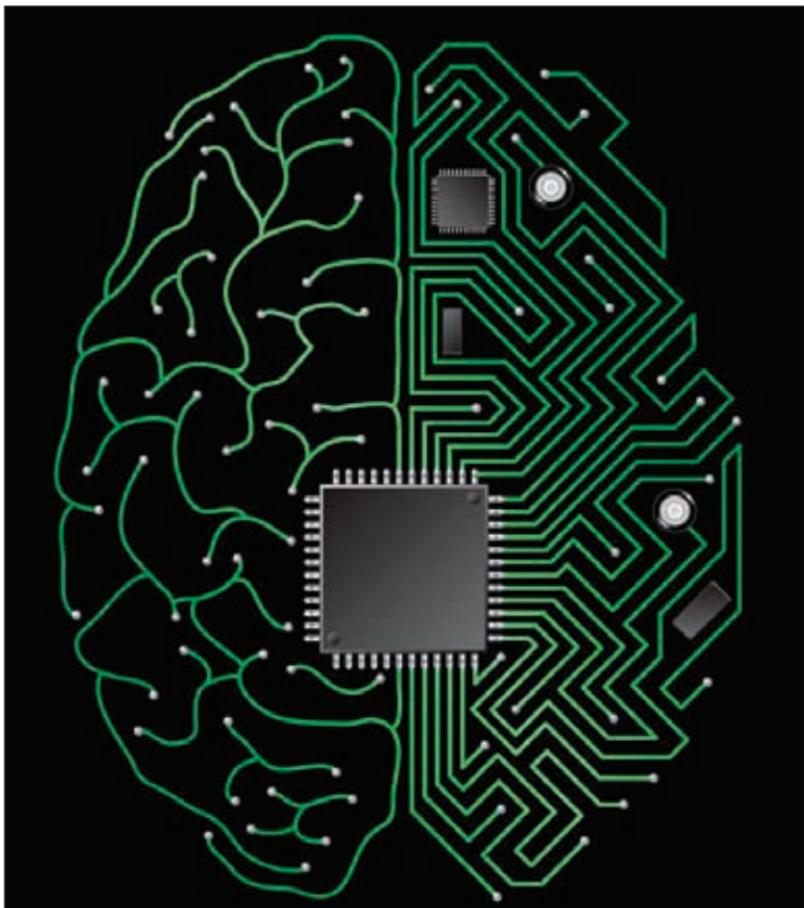
Artificial intelligence has already gotten to the point where the kinds of things computers can do are jaw-dropping. Yet we are still at an early stage and there's a long way to go. What will the next wave of scientists work on and how long will it take for the dreams of science fiction writers to be fulfilled in real life?

The question of time needed for AI to mature has been the most oft-asked and oft-debated among cognitive scientists. In fact, just like chess grandmaster David Levy made the famous bet that no computer would beat

him for 10 years, there are a lot of other such “long bets” about intelligent agents. The most famous of them is the one between the founder of Lotus, Mitch Kapor and futurist Ray Kurzweil.

Kapor contends that “By 2029 no computer – or ‘machine intelligence’ – will have passed the Turing Test.” Kurzweil argued against this prediction, with the two famously laying a bet of \$20,000 that will be administered by the Long Now Foundation, with the proceeds going to charity.

But Kurzweil is quick to add a cautionary note about the implications of a machine being able to pass the Turing Test. He says that once a machine does pass the Turing Test to show human-like intelligence, it will necessarily soar past it after that.



Recreating the brain

“Electronic circuits are already at least 10 million times faster than the electrochemical information processing in our interneuronal connections. Machines can share knowledge instantly, whereas we biological humans do not have quick downloading ports,” he says.

Thus “Turing Test level” machines will be able to combine human level intelligence with the powerful ways in which machines already excel. In addition, machines will continue to grow exponentially in their capacity and knowledge. It will be a formidable combination.

Recreating the brain

Every conversation about true artificial intelligence will go round and round and wind up at the same point: are we able to fully recreate the human brain? Asimov’s vision of the positronic brain might be a fair distance away, but a group of researchers has made rapid strides in recreating the human brain.

At the Ecole Polytechnique Fédérale de Lausanne (EPFL) research institute in Switzerland, scientists are working to create a simulation of the human brain on a supercomputer. The idea of the Blue Brain project is to reverse-engineer how the human brain works.

The first step the scientists took was study the brain of a rat, since rat brains have been shown to be the most biologically similar to humans in the mammalian world. The crux of the project was a virtual revisualisation of the neocortical column – the most evolved part of the brain which is responsible for diverse tasks, the primary among them being reasoning. Although the human brain has a lot more neocortical columns than a rat, the researchers simply wanted to create one column, with the notion that if they can make one, they can make many more.

In 2006, the project tasted success and drew high praise from all quarters. The team is now working on remodelling the entire rat brain in its simulation on a super-computer, after which they will look at doing the same with a cat’s brain (more complex), a monkey’s brain (even more complex) and then finally a human brain.

But the scientists are already working alongside to turn their virtual representation into a real, physical one. After a lot of analysis and trial-and-error, the team finally decided that carbon nanotubes were the way forward in creating a physical circuit capable of replicating the neurons of the brain. Like neurons, carbon nanotubes are highly electrically conductive, forming extremely tight contacts with neuronal cell membranes. Unlike the metal

electrodes that are currently used in research and clinical applications, the nanotubes can create shortcuts between the two extremes of a neuron, resulting in enhanced neuronal excitability.

For Professor Henry Markram, project leader, the most exciting part of his research is putting together the hundreds of thousands of small pieces of data that his lab has collected over the past 15 years, and seeing what a microcircuit of the brain looks like.

“When we first switched it on, it already started to display some interesting emergent properties,” he says. “But this is just the beginning because we know now that it is possible to build it. As we progress we are learning about design secrets of our brains which were unimaginable before. In fact the brain uses some simple rules to solve highly complex problems and extracting each of these rules one by one is very exciting. For example we have been surprised at finding simple design principles that allow billions of neurons to connect to each other. I think we will understand how the brain is designed and works before we have finished building it.”

Meanwhile, a different group of scientists at the California Institute of Technology has successfully made an artificial neural network out of DNA, creating a circuit of interacting molecules that can recall memories based on incomplete patterns, just as a human brain can.

While this proof-of-principle experiment shows the promise of creating DNA-based networks that can—in essence—think, this neural network is limited, the researchers say. The human brain consists of 100 billion neurons, but creating a network with just 40 of these DNA-based neurons – 10 times larger than the demonstrated network – would be a challenge, according to the researchers.

Markram, however, is optimistic and believes that the human brain could be replicated in the next 10 years: “I absolutely believe it is technically and biologically possible.”

“The brain is, of course, extremely complex because it has trillions of synapses, billions of neurons, millions of proteins, and thousands of genes. But they are still finite in number. Today’s technology is already highly sophisticated and it allows us to reverse engineer the brain rapidly.”

Love and sex

According to David Levy – yes, the same chess grandmaster who had once predicted that no computer would beat him – we’re fast approaching an age where it would be possible for humans to fall in love with robots

and have sex. In fact, he predicts that this should be a likely scenario by the year 2050.

Levy, who has a Ph.D. on the subject of human-robot relationships, points out how when the chatterbot ELIZA was first invented, some test subjects reported being unusually attracted to it. He found similar evidence of some individuals preferring to have “relationships” with their computers than with other humans.

“At first, sex with robots might be considered geeky,” Levy told LiveScience magazine. “but once you have a story like ‘I had sex with a robot, and it was great!’ appear some place like Cosmo magazine, I’d expect many people to jump on the bandwagon.”

The reason Levy is so confident about the gradual transition is because of his understanding of both – artificial intelligence and love. In his thesis *Intimate Relationships with Artificial Partners*, Levy writes that several of the necessary requirements for humans to fall in love are already present in robots and it’s only a matter of time before they fully encompass these.

From the artificial intelligence side, Levy says that robots are already becoming more and more humanoid in appearance and behaviour. Within a



Wall-e and Eva in love is cute, but...



Illustration: Rudy-Jan Eeber (rudyjabeer.com)

...would society find this type of love "cute"?

few decades, they will be even more humanlike, expressing emotions and possessing a personality, social instincts and physical (albeit artificial) features like pheromones – all necessary elements needed to fall in love. “Many people will be falling in love with them, having sex with them, and even marrying them,” Levy writes.

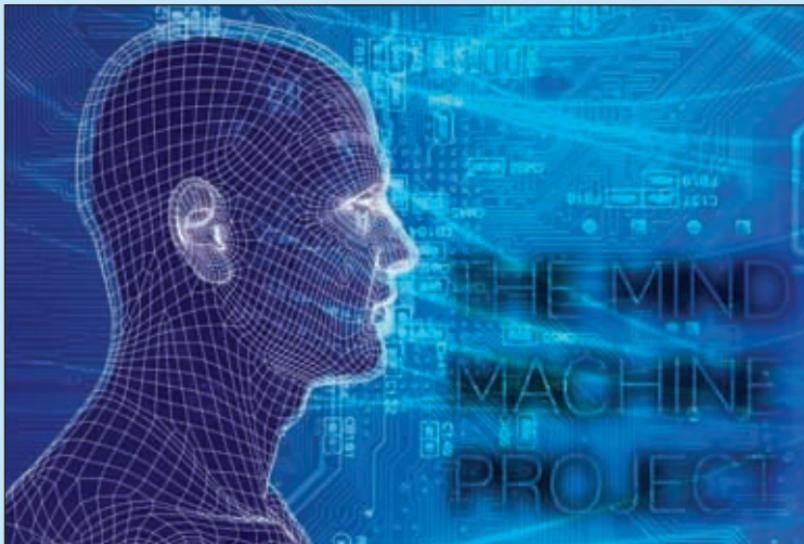
On the human side, he contends that trends in mankind’s history are such that it seems like

a natural progression.

When the virtual pet Tamagotchi or the AIBO robotic dog first made an appearance, they were ridiculed. But slowly, people grew to appreciate them and form a bond as closely-knit as with a real pet.

Levy also observes that as society progresses, our attitude towards sex and sexual orientation starts becoming more liberal.

Of course, all of this throws up a plethora of ethical questions, but as Levy points out, it’s not about looking at things as a whole but more about looking at things from the perspective of a human relationship. In an interview with the *Scientific American*, he says, “It’s not that people will fall in love with an algorithm, but that people will fall in love with a convincing simulation of a human being, and convincing simulations can have a remarkable effect on people.” **d**



LOOKING AT AI UNDER A NEW LIGHT

While artificial intelligence has been a form of science for several decades, a group of boffins now wants to rethink the very way we look at this field

In the previous chapter, we saw that computer scientists and futurists like Markram and Levy are optimistic about where artificial intelligence is headed. But it could be argued that their ideas remain firmly grounded in the way the field was defined by the thinkers in the 1950s and 60s, like Herbert Simon, John McCarthy and Roger Schank.

Some scientists think that this approach is archaic and needs to be rethought given what we know about the human mind and machinery now. Spearheaded by the Massachusetts Institute of Technology, some of the pioneers of the AI field have teamed up with new generations of thinkers to perform a massive “do-over” of the whole idea.

This time, they’re determined to get it right – and, with the advantages of hindsight, experience, the rapid growth of new technologies and insights from the new field of computational neuroscience, they think they have a good shot at it.

The new project is called the “Mind Machine Project” or “MMP”, a loosely bound collaboration of about two dozen professors, researchers, students and postdocs. It’s led by Newton Howard, who came to MIT to head this project from a background in Government and Industry Computer Research and Cognitive Science.

According to Neil Gershenfeld, one of the leaders of MMP and director of MIT’s Center for Bits and Atoms, one of the project’s goals is to create intelligent machines – “whatever that means.”

The project is “revisiting fundamental assumptions” in all of the areas encompassed by the field of AI, including the nature of the mind and of memory, and how intelligence can be manifested in physical form, says Gershenfeld, professor of media arts and sciences. “Essentially, we want to rewind to 30 years ago and revisit some ideas that had gotten frozen,” he says, adding that the new group hopes to correct “fundamental mistakes” made in AI research over the years.

The birth of AI as a concept and a field of study is generally dated to a conference in the summer of 1956, where the idea took off with projections of swift success. One of that meeting’s participants, Simon, predicted in the 1960s, “Machines will be capable, within 20 years, of doing any work a man can do.” Yet two decades beyond that horizon, that goal now seems to many to be as elusive as ever. It is widely accepted that AI has failed to realise many of those lofty early promises.

Fixing what’s broken

Gershenfeld says he and his fellow MMP members “want to go back and fix what’s broken in the foundations of information technology.” He says that there are three specific areas – having to do with the mind, memory, and the body – where AI research has become stuck, and each of these will be addressed in specific ways by the new project.

The first of these areas, he says, is the nature of the mind: “How do you model thought?” In AI research to date, he says, “what’s been missing is an ecology of models, a system that can solve problems in many ways”, as the mind does.

Part of this difficulty comes from the very nature of the human mind, evolved over billions of years as a complex mix of different functions and systems. “The pieces are very disparate; they’re not necessarily built in a compatible way,” Gershenfeld says. “There’s a similar pattern in AI research. There are lots of pieces that work well to solve some particular problem, and people have tried to fit everything into one of these.” Instead, he says, what’s needed are ways to “make systems made up of lots of pieces” that work together like the different elements of the mind. “Instead of searching for silver bullets, we’re looking at a range of models, trying to integrate them and aggregate them,” he says.

The second area of focus is memory. Much work in AI has tried to impose an artificial consistency of systems and rules on the messy, complex nature of human thought and memory. It’s now possible to accumulate the whole life experience of a person, and then reason using these data sets which are full of ambiguities and inconsistencies. That’s how the human brain functions – we don’t reason with precise truths. Computers need to learn ways to reason that work with, rather than avoid, ambiguity and inconsistency.

And the third focus of the new research has to do with what the MMP describes as “body”: “Computer science and physical science diverged decades ago,” Gershenfeld says. Computers are programmed by writing a sequence of lines of code, but the human mind doesn’t work that way. In the mind, everything happens everywhere all the time. A new approach to programming, called RALA (Reconfigurable Asynchronous Logic Automata) attempts to re-implement all of computer science on a base that looks like physics. The computations have physical units of time and space, so the description of the system aligns with the system it represents. This could lead to making computers that run with the fine-grained parallelism the brain uses.

Rethinking the Turing Test

Marvin Minsky, one of the pioneering researchers from AI’s early days and part of the MMP team, sees real hope for important contributions this time around. The inability of modern AI to pass the Turing Test is partly to blame on the way we assign credence to the human mind as a benchmark for intelligence, he says.

Now, Minsky proposes a different test that would determine when machines have reached a level of sophistication that could begin to be truly useful. His method? To test whether the machine can read a simple children's book, understand what the story is about, and explain it in its own words or ask reasonable questions about it.

It's not clear whether that's an achievable goal on this kind of timescale, but Gershenfeld says, "We need good challenging projects that force us to bring our program together."

A co-processor for the human brain

Of course, as with all talks about artificial intelligence, there comes a time when you have to breach the topic of hybrid human-artificial minds, or "bionic brains". While the scientists don't want to speculate about that, they are working on something that they like to call a "co-processor for the brain".

This system, also referred to as a cognitive assistive system, would initially be aimed at people suffering from cognitive disorders such as Alzheimer's disease. The concept is that it would monitor people's activities and brain functions, determine when they needed help, and provide exactly the right bit of helpful information at the right time. For example, with a faulty memory in the patient, the co-processor could tell the brain the name of a person who just entered the room, and information about when the patient last saw that person.

The same kind of system could also find applications for people without any disability, as a form of brain augmentation – a way to enhance their own abilities, for example by making everything from personal databases of information to all the resources of the internet instantly available just when it's needed. The idea is to make the device as non-invasive and unobtrusive as possible – perhaps something people would simply slip on like a pair of headphones.

In the longer run, Minsky still sees hope for far grander goals. For example, he points to the fact that his iPhone can now download thousands of different applications, instantly allowing it to perform new functions. Why not do the same with the brain? "I would like to be able to download the ability to juggle," he says. "There's nothing more boring than learning to juggle." Is the Matrix real? 

Control Panel > System and Security > System

thinkdigit.com/review.php

EENIE MEENIE
MYNIE MO

NEWS AND REVIEWS

Comprehensive news, unbiased reviews and ratings to help you make the right purchase



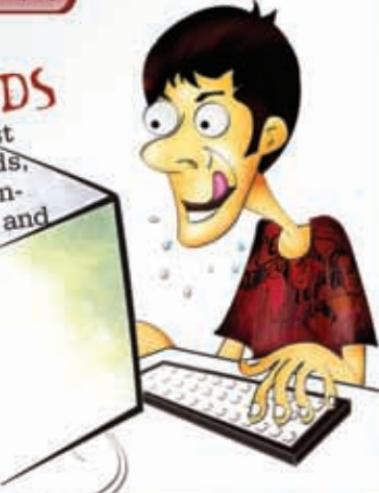
camera, laptop, mobile phones, digital cameras

Control Panel > System and Security > System

thinkdigit.com/download.php

DOWNLOADS

Check out the latest software downloads, games for your Windows, Linux, Mac and PDA/mobiles



All this and more in the world of Technology

**VISIT
NOW**



www.thinkdigit.com

Join 60000+ members of the Digit community



<http://www.facebook.com/thinkdigit>



Your favourite magazine on your social network. Interact with thousands of fellow Digit readers.

<http://www.facebook.com/IThinkGadgets>



An active community for those of you who love mobiles, laptops, cameras and other gadgets. Learn and share more on technology.

<http://www.facebook.com/Gladtobeaprogrammer>



If you enjoy writing code, this community is for you. Be a part and find your way through development.

<http://www.facebook.com/ithinkgreen>



Make a difference to your world and your planet. Join our community and show that you care.